

Contents

7	BTeV R&D Program	1
7.1	Pixel Detector R&D	2
7.1.1	Introduction	2
7.1.2	Sensor development	2
7.1.3	Pixel readout chip	5
7.1.4	Bump bonding development	7
7.1.5	Multichip Module	8
7.1.6	Mechanical support	9
7.1.7	Future R&D Plan	12
7.2	Silicon Forward Tracker R&D	16
7.2.1	Silicon microstrip sensors	16
7.2.2	Front-end read-out electronics	17
7.2.3	Detector support mechanics and cooling	18
7.2.4	DAQ system	19
7.3	Straw Tube Tracker R&D	25
7.3.1	Single Straw Prototype	25
7.3.2	Wire Stringing and Tension Measurement	26
7.3.3	Straw Stretch/Tension Tests	27
7.3.4	Two-Module Prototype Construction	27
7.3.5	Other R&D	29
7.3.6	Simulations	30
7.4	Ring Imaging Cherenkov Detector R&D	32
7.4.1	HPD Development	32
7.4.2	HPD Front End Readout	32
7.4.3	HPD High Voltage	34
7.4.4	Mirrors	35
7.4.5	Liquid Radiator	36
7.4.6	Magnetic Shielding of Photodetectors	39
7.4.7	Radiation Damage Studies	39
7.4.8	Mechanical Design	40
7.4.9	Beam Test	40
7.5	Electromagnetic Calorimeter R&D	42

7.5.1	Introduction	42
7.5.2	PbWO ₄ Crystals	42
7.5.3	Mechanical Support Structure	46
7.5.4	Electronics	48
7.6	BTeV Muon System R&D	49
7.6.1	Introduction	49
7.6.2	Summary of the BTeV muon system	49
7.6.3	Past research and development work	51
7.6.4	Future research and development work	54
7.7	Trigger R&D	55
7.7.1	Level 1 Vertex Trigger	55
7.7.2	Level 1 Muon Trigger	58
7.7.3	Level 1 Hardware	59
7.7.4	Level 2	61
7.7.5	Level 3	63
7.8	Data Acquisition R&D	66
7.8.1	Overview	66
7.9	Data Acquisition Hardware	66
7.9.1	Data Acquisition Software	67
7.9.2	Summary and Future Plans	68
7.10	BTeV Real-Time Embedded Systems R&D	69
7.10.1	Introduction	69
7.10.2	Project Overview	69
7.10.3	Strong Connection to BTeV	72
7.10.4	Group Members	73
7.10.5	Current and Proposed Activities	73

Chapter 7

BTeV R&D Program

Introduction

In this chapter, we describe the status of the BTeV R&D program, including its recent accomplishments and its plans for the next year, or, in some cases, two years.

R&D is critical to the ultimate success of BTeV. We use established technologies where possible. However, in order to meet the requirements described in Chapter 3, several new technologies have to be developed. In some cases, even where we have chosen “existing technologies,” they are still relatively new technologies which have been developed by others, but have not yet been tested in large scale operation. In some cases where we have chosen existing technologies, we will use them in different ways or expose them to different environments than the ones they were developed for. This means that we must do some research and development, or at least some development, on every subsystem.

In the following, we describe the R&D effort associated with each major detector component, the trigger, and the data acquisition system. The various activities include development of analog and digital ASICs, studies of radiation hardness, ageing studies, simulations, finite element analyses, materials testing, mechanical mockups, bench tests, beam tests, reliability studies, and production yield studies, to name a few. A prototype will be constructed for every detector subsystem and for the trigger and data acquisition system. All detector prototypes will be operated in test beams over the next two years. All materials will be tested for radiation hardness at the required level and for ageing and environmental robustness. Critical calibration systems will also be tested thoroughly. It is through these efforts that we intend to guarantee that our designs will meet our requirements and that we can build the BTeV detector on time and on budget.

The R&D program has been supported by DOE through Fermilab and the DOE University program, NSF, INFN, and IHEP. We have also received significant support from the universities whose physicists are participating in the experiments.

7.1 Pixel Detector R&D

7.1.1 Introduction

One of the defining characteristics of BTeV is the plan to use a vertex trigger as the primary trigger for the experiment. The pixel detector is required to provide high resolution space points (better than $9\mu\text{m}$ for tracks of all angles) which will be used by this trigger. This imposes both fast readout and large bandwidth requirements on the front end electronics. To minimize the extrapolation error, the detector will be placed as close as possible to the interaction point and hence will be exposed to a significant level of irradiation. At our maximum projected luminosity, it is expected that the innermost pixel detector will receive a fluence of 1×10^{14} minimum ionizing particles/cm²/year. This significant radiation environment means that all components of the pixel system have to be radiation hardened.

Since the submission of the BTeV Proposal in the spring of 2000, we have made great progress in the development of the individual components required to build the BTeV pixel detector. Our R&D has also started to address the system engineering aspects. This section describes the main accomplishments achieved during this period.

The major components of the pixel detector system are the sensor, readout chip, sensor-readout-chip connection (bump bonding), high-density interconnection between the pixel readout chips and the system control elements, and the mechanical support and cooling system. We have been designing and purchasing these components, assembling units and testing them in beams and exposing them to intense radiation. We have also performed detailed simulation studies to understand the various design issues for the components as well as system aspects. Through these efforts, not only are we learning what is needed for BTeV, but we are gaining the necessary experience and know-how to build the actual pixel detector for the BTeV experiment. One of the highlights of this effort is the successful demonstration in a test beam during the 1999 Fermilab fixed target run of the resolution that can be achieved with a pixel detector. Fig. 7.1 shows the resolution as a function of the incident beam angle for a pixel detector [1]. Two curves and data points are included in the figure: the solid line and circles show prediction and measurements done with an eight-bit ADC external to the pixel readout chip to take charge sharing between pixels into account; the dashed curve and triangular data points illustrate the simulation and measurements obtained if we were to use only binary (on-off) readout. The clear advantage of using charge sharing, made possible by analog readout, is evident and for all incident angles, a resolution of better than $9\mu\text{m}$ has been obtained.

7.1.2 Sensor development

The main challenge is to have a radiation hardened detector which will survive and remain operational after significant radiation damage to both the surface and the bulk of the silicon sensors.

The bulk damage is mainly due to the non-ionizing energy loss (NIEL) which, through the displacement of atoms in the crystal lattice, creates new energy levels, effectively acting

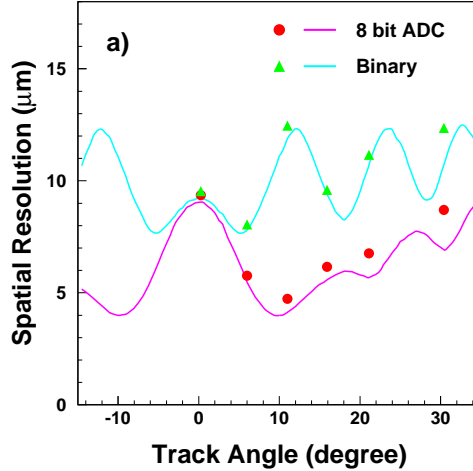


Figure 7.1: Resolution as a function of the angle of the incident beam. Data was taken with prototype pixel detectors during the 1999 Fermilab fixed target run. The detectors were instrumented with the earliest versions of the pixel readout chip FPIX0 at Fermilab. The curves represent the predicted resolution: the oscillating curve is the simulated digital resolution and the lower one assumes 8-bit charge digitization. The circles and triangles are extracted from the data.

as acceptors. Therefore the effective doping concentration will change with irradiation. For very-high-dosage irradiation, this will eventually lead to inversion of the conduction type of the bulk material (type-inversion), increases in leakage current and depletion voltage, changes in capacitance and resistivity, and charge collection losses. These are problems that need to be addressed by all the next generation hadron collider experiments. As a result, there is a worldwide effort to address these technical challenges. Solutions include the design of multiple guard ring structures to avoid avalanche breakdown along the sensor edges, low resistivity silicon substrates to delay type inversion, thin detectors to reduce the depletion voltage required, and oxygenated silicon wafers to reduce the effects of radiation-induced defects in the silicon lattice [2].

In order to increase the useful operation time of the silicon sensors, operation with partial depletion has to be considered. This is more suitable for n-type pixel readout, because after type inversion the depleted region will grow from the n+ side of the junction. For this reason, the BTeV pixel sensors have $n^+/n/p^+$ configuration. In these detectors, the charge collecting pixels are defined by the n-implants that are isolated from their neighbors. Without isolation, the accumulation layer induced by the oxide charge would short the individual n^+ pixels together. We have explored two isolation techniques:

- The p-stop isolation where a high dose p-implant surrounds the n-region.

- The p-spray isolation developed by the ATLAS collaboration, where a medium dose shallow p-implant is applied to the whole n-side. To increase the radiation hardness and also the breakdown voltage before irradiation, a “grading” of the p-spray implantation (moderated p-spray) is required [3]

We have tested prototype p-stop sensors produced by SINTEF. The base material is low resistivity (1.0-1.5 K Ω cm) silicon, in a $\langle 100 \rangle$ lattice, 270 μm thick. Some of the wafers have been oxygenated. A few of these sensors have been exposed to a 200 MeV proton beam at the Indiana Cyclotron Facility (IUCF). To characterize these sensors before and after irradiation, we measured bulk parameters of the sensors including the bias voltage dependence of the leakage current, the full depletion voltage, breakdown voltage, and the temperature dependence of the leakage current [4].

Figure 7.2 shows the leakage current measurements before and after irradiation up to a fluence of 2×10^{14} 200 MeV protons cm^{-2} for a SINTEF p-stop sensor. The measurement after irradiation was performed at various temperatures (23°C, 10°C, 0°C, and -10°C) and as expected, we observed that the leakage current decreases exponentially with temperature.

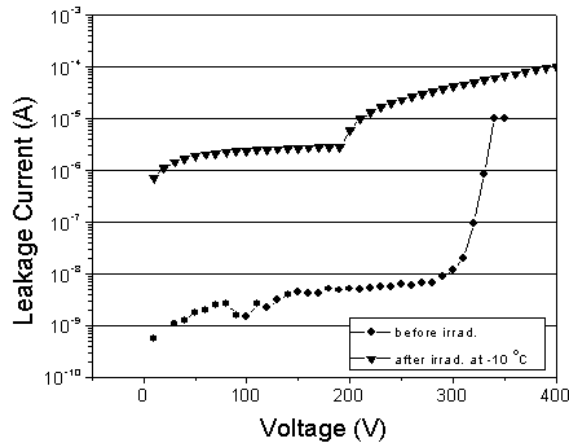


Figure 7.2: Leakage current measurements before (at room temperature) and after (at -10°C) irradiation to 4×10^{14} p/cm 2 for a SINTEF p-stop sensor.

The other bulk damage is the change in effective doping density which is reflected in a change in the full depletion voltage. Fig. 7.3 shows the dependence of the full depletion voltage on the proton irradiation fluence for a few p-stop sensors made from standard and oxygenated wafers. At a fluence of 4×10^{14} p cm^{-2} , the full depletion voltage is still rather low, even lower than the value before irradiation. This characteristic is due to the low resistivity of the starting silicon material. This result, together with the fact that the breakdown voltage is still high compared to the full depletion voltage after irradiation, means that the BTeV pixel detector can be fully depleted without excessively high bias voltage even after a few years of operation.

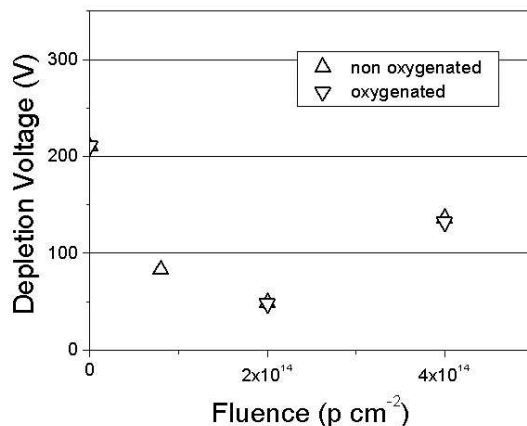


Figure 7.3: Full depletion voltage as a function of the fluences of the proton irradiation for normal and oxygenated sensors.

We have recently also started to characterize moderated p-spray sensors obtained from three vendors via the ATLAS collaboration. Typically, the initial leakage current of these detectors was found to be about 10 nA/cm², similar to the p-stop sensors. Furthermore, the breakdown voltage before irradiation, for these sensors, is rather high as expected. During the next few months, we will perform irradiation studies on these sensors.

We also plan to study both the p-stop and the moderated p-spray detectors in a test beam to study the charge collection properties before and after irradiation and compare the results with the predictions from simulation.

7.1.3 Pixel readout chip

The use of the pixel detector data in the first level trigger means that the BTeV pixel readout chip must be capable of reading out all hit information from every $p\bar{p}$ interaction. Furthermore, the pixel readout chip should be optimized for the 132 ns time between crossings planned for the Tevatron collider. It must be radiation hard so that it can be used close to the beamline. This requires a pixel readout chip with a low noise front-end, an unusually high output bandwidth, and implemented in a radiation-hard technology. During the last few years, a pixel readout chip has been developed at Fermilab to meet these requirements. This has been done through several stages of chip development, each of increasing complexity [5].

Recently, the pixel readout chips have been implemented in two different commercial 0.25 μ m CMOS processes following radiation tolerant design rules (enclosed geometry transistors and guard rings) [6]. The preFPIX2I chip, containing 16 columns with 32 rows of pixel cells, and complete core readout architecture, was manufactured by a vendor through CERN [7]. The preFPIX2Tb chip, contains, in addition to the preFPIX2I chip features, a new programming interface and 14 digital-analog-converters (DAC) to control the operating

and threshold settings of the whole chip. It was manufactured by Taiwan Semiconductor Manufacturing Company (TSMC).

To study total dose and Single Event Effects (SEE), samples of these prototype chips have been exposed to 200 MeV protons at IUCF. The comparison of the chip performance before and after exposure shows the high radiation tolerance of the design [8]. Chips have been exposed to as much as 7.4×10^{14} protons-cm⁻² (about 43 Mrad) and no evidence of catastrophic failure or deterioration of the functionality of the readout chip has been observed. In particular, no radiation induced SEE, such as Latch-Up or Gate-Rupture has been observed. Fig.7.4 shows the effect of radiation on amplifier noise.

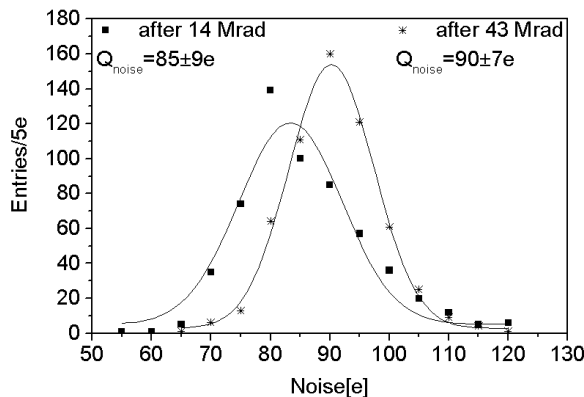


Figure 7.4: Measured noise in the 576 pixel cells of preFPIX2Tb chip after 14 and 43 Mrad of proton irradiation.

An important feature of the preFPIX2Tb chip is the implementation of on-chip DAC's in order to minimize the number of external I/O lines. The change of the DAC behavior due to the proton irradiation has been measured and is shown in Fig.7.5. The three curves shown correspond to the deviation from the linear fit to the unirradiated data for total dose of 0, 14, and 43 Mrad. It can be seen that the linearity and accuracy of the DAC output remains acceptable after 43 Mrad total dose.

We have also measured the Single Event Upset (SEU) cross-section of static registers implemented on the preFPIX2Tb chips, in order to establish the sensitivity of our design to radiation induced digital soft errors. The measurements consisted of detecting bit error rates in the static registers controlling the readout chip front-end operating conditions and the pixel cell response. The single bit upset cross-section measured for the DAC's located on the chip periphery was $(5.5 \pm 0.6 \pm 0.5) \times 10^{-16}$ cm² while for the mask and charge-injection registers located inside each pixel cell was $(1.9 \pm 0.2 \pm 0.2) \times 10^{-16}$ cm² (where the first error is statistical and the second systematic due to uncertainty in the beam fluence) [9]. We tested and did not observe any dependence of the upset rate on the beam incidence angle or clock frequency up to 16 MHz.

Our measurements of the SEU rate implies that the SEU bit error rate in the BTeV pixel detector is small enough that it will not be necessary to design explicitly SEU tolerant

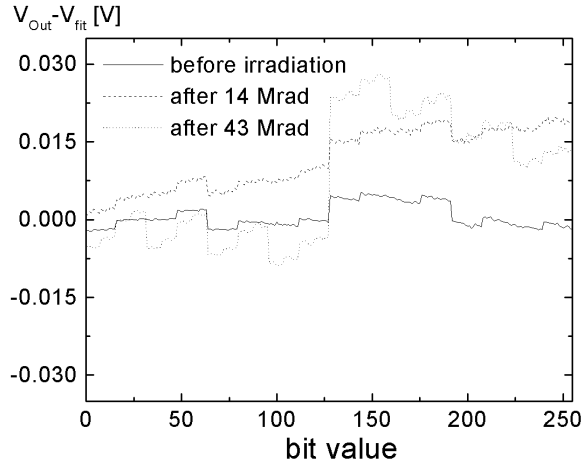


Figure 7.5: DAC analog response before and after 14 and 43 Mrad total dose exposure to 200 MeV protons. The full scale (255 counts) corresponds to about 1.7V

registers. Rather, the SEU rate can be comfortably handled by a periodic readback of the chip configurations during data-taking and a download of the chip configuration whenever an upset is detected.

Based on the experience gained, we intend to submit a full-size BTeV pixel readout chip soon. That chip will have 22 columns by 128 rows. The new section to be included in the chip is a high speed data output interface which accepts data from the core, serializes the data, and transmits the data off chip.

7.1.4 Bump bonding development

The BTeV pixel detector, like all other pixel systems used in or planned for HEP experiments, is based on a hybrid design. With this approach, the readout chip and the sensor array are developed separately and the detector is constructed by flip-chip mating the two together. This method offers maximum flexibility in the development process, choice of fabrication technologies, and sensor materials. However, it requires the availability of a highly reliable, reasonably low cost fine-pitch flip-chip attachment technology. We have focused our study on two options: indium bumps, and Pb-Sn solder bumps.

A series of yield and stability tests were performed on bump-bonded test structures. These tests were done with indium, fluxed-solder, and fluxless-solder bumps from a number of commercial vendors. Our tests have validated the use of indium and fluxless-solder as viable technologies. The failure rate obtained from this large scale test is about 2×10^{-4} which is adequate for our needs [10].

In order to check the long term reliability of the bump-bonding technology, we monitored the quality of the connectivity over a period of one year. In addition, we performed thermal cycling: exposure to -10°C for 144 hours and $+90^{\circ}\text{C}$ for 48 hours in vacuum. Furthermore,

we irradiated some of these test structures with a ^{137}Cs source up to a dose of 13 Mrad. The typical failure rate of both types of bumps under these stringent tests was found to be a few $\times 10^{-4}$. These results show that both techniques are highly reliable [11].

One of the remaining concerns is that thermal stress on the bumps due to the Coefficient of Thermal Expansion (CTE) mismatch of the bump material, silicon, and the substrate material on which the detector is placed. We have recently embarked on a study of this problem and results are expected in the next few months.

7.1.5 Multichip Module

Each pixel readout chip includes a high density of control and data output lines at the periphery. These lines need to be connected to the back-end electronics. A full set of pads is available on the readout chip for these interconnection purposes. This is achieved through a high density, low mass flex circuit wire bonded to a number of readout chips to form a multichip module.

Each pixel half-plane will be made up of a number of these multichip modules. The module is the basic building block of the pixel detector system.

The pixel module is composed of three layers. The lowest layer is formed by the readout integrated circuits (ICs). The back of the ICs is in thermal contact with the supporting structure, while the top is flip-chip bump-bonded to the pixel sensor. A low mass flex-circuit interconnect is glued on the top of this assembly, and the readout IC pads are wire-bonded to the flex-circuit (see Fig.7.6).

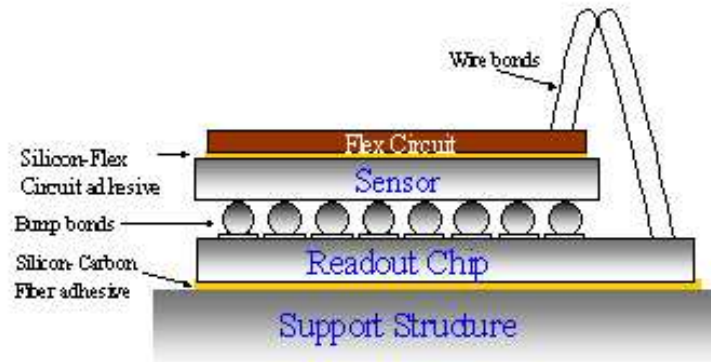


Figure 7.6: Sketch of the pixel multichip module stack

We are presently developing prototypes of the pixel module based on the earlier version of the pixel readout chip FPIX1 to assess the electrical and mechanical performance of such assembly, as well as to acquire early insights on the construction process and yield. A previous module was already manufactured with great success [12].

The FPIX1 interface with the data acquisition system was not optimized to reduce the number of interconnections. The large number of signals in this prototype imposes space

constraints and requires aggressive circuit design rules, such as $35\mu\text{m}$ trace width and trace-to-trace clearance of $35\mu\text{m}$ and four metal layers. A circuit with such characteristics is very difficult to obtain and very few places have such manufacturing expertise. The Engineering Support and Technical Division at CERN manufactured the FPIX1 interconnect flex circuit. Fig. 7.7 shows a picture of the flex circuit. Several design strategies to minimize electrical noise and guarantee signal integrity were incorporated in the layout and are being evaluated.

The interface adhesive between the flex-circuit and the pixel sensor has to compensate for mechanical stress due to the CTE mismatch between the flex circuit and the silicon pixel sensor. Two alternatives are being pursued.

Two partial assemblies of this pixel module have already been characterized. One is a single FPIX1 readout chip flip-chip mated to a SINTEF sensor using indium bump bonds. The flex circuit is pasted on top of the sensor. The second pixel module has three readout chips wire-bonded but does not have a pixel sensor.

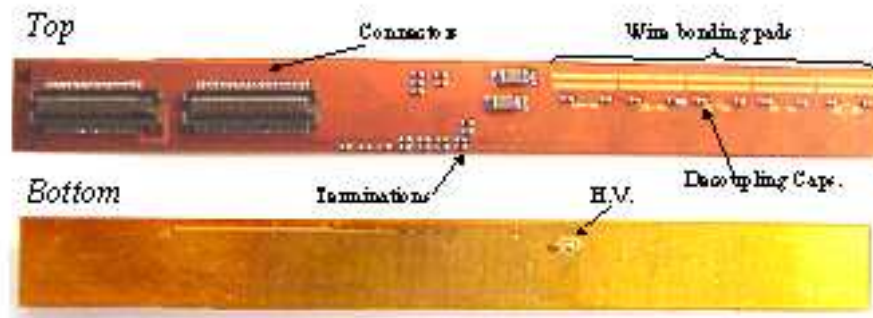


Figure 7.7: Picture of the flex-circuit made by CERN.

These modules were characterized for noise and threshold dispersion. These characteristics were measured by injecting test charge in the analog front end of the readout chip with a pulse generator. The results for various thresholds are summarized in Table 7.1 and 7.2 [13]. These results are comparable with previous characterization results of single readout IC mounted on a printed circuit board. No crosstalk problem has been observed between the digital and analog sections of the readout chip and the flex circuit.

7.1.6 Mechanical support

Significant progress has been made on the engineering design of the overall mechanical support, the vacuum vessel, motor drive assembly, and the individual substrates on which the pixel modules will be mounted. In some cases, early prototypes have been made and evaluated.

Each pixel half-station is assembled on two substrates, with the pixel modules placed with a small overlap on both surfaces of the substrate to provide complete coverage of the active area. For a number of years, the baseline design was to use a substrate made out of a novel

<i>Single bare chip</i>				<i>Single chip with sensor</i>			
μ_{Th}	σ_{Th}	μ_{Noise}	σ_{Noise}	μ_{Th}	σ_{Th}	μ_{Noise}	σ_{Noise}
7400	360	75	7	7800	410	94	7.5
6400	330	78	12	6500	390	110	11
5400	390	79	11	5500	380	110	13
4400	380	78	11	4400	380	110	15
3500	380	79	12	3350	390	120	20
2500	370	77	13	2300	390	120	21

Table 7.1: Performance of the one-chip FPIX1 module without and with sensor. μ represents mean values of distributions and σ 's are their RMS. All numbers are given in equivalent electrons. There is no significant increase in noise and threshold dispersion with the sensor attached.

<i>Setting</i>	<i>Chip 1</i>				<i>Chip 2</i>				<i>Chip 3</i>			
$V_{Th}(V)$	μ_{Th}	σ_{Th}	μ_{Noise}	σ_{Noise}	μ_{Th}	σ_{Th}	μ_{Noise}	σ_{Noise}	μ_{Th}	σ_{Th}	μ_{Noise}	σ_{Noise}
1.95	5210	451	93	11	5100	501	105	12	5900	432	85	10.5
2.05	3305	412	80	10	3340	390	92	12	4200	410	87	11
2.15	1540	440	77	11	1850	420	79	11				

Table 7.2: Performance of the three-chip FPIX1 module. μ represents mean values of distributions and σ 's are their RMS. All numbers are given in equivalent electrons.

material called "fuzzy carbon" with a number of embedded cooling tubes made out of glassy carbon. However, fuzzy carbon is very fragile and is made by a proprietary process owned by a single vendor. More importantly, such a design will have a large number of cooling joints and pipes containing coolants placed inside a vacuum system. The reliability and the risk of a leak in the system is a subject of grave concern. On another front, the outgassing tests of a 5 % model of the pixel detector at various temperatures suggested that the use of a cryogenic panel at -160°C might provide sufficient pumping to achieve the required vacuum level. The presence of the cryogenic panels and liquid nitrogen lines inside the pixel vacuum vessel provides a convenient heat sink. Cooling for the pixel substrate can now be done by conduction without the need of flowing coolant through the substrates. We will then have a joint-free and potentially leak-tight cooling system. A material with very high thermal conductivity is needed for this kind of heat transfer mechanism in order to minimize the temperature gradient across the substrate. After some preliminary study among carbon-carbon, carbon-fiber reinforced plastics, carbon foam, flexible pyrolytic graphite sheet (PGS), and thermal pyrolytic graphite (TPG), TPG was chosen because of its outstanding thermal properties and low radiation length. To provide relief caused by the stresses due to the difference in CTE amongst the various materials that will be used (e.g. TPG, carbon fiber

support, cooling blocks), the more flexible and light weight PGS will be used to connect the TPG substrate to the cooling blocks. The thermal conductivity of TPG, can be as high as 1,700 W/m-C at room temperature. This property is temperature dependent, and it even surges to a peak of about 3,000 W/m-C at -160°C. TPG is currently used by the ATLAS SCT barrel modules and outer forward silicon modules. It has also been used by HERA-B and AMS and is proven to be a good candidate for such substrate design.

Figure 7.8 shows the design of the vacuum vessel and support structure for the pixel detector. Each pixel station will be attached by brackets to a C-shaped support frame. Both the bracket and the frame will be made out of carbon fiber material. Signal feedthrough the vacuum vessel is done by a number of multilayer printed circuit boards with high density multipin connectors. Actuator assemblies located outside the vacuum vessel will be used to control the movement of the pixel stations in and out of the beam during data-taking and beam refill.

The major assembly steps have been worked out for the current design. The mechanical strength of all the important elements such as substrate mounting brackets, C-fiber support cylinder, and vacuum vessel were checked with finite-element-analysis (FEA) calculations to make sure that any deflections and stresses under load are acceptable. To verify the calculations, the manufacturing process, and assembly procedure, a prototype support cylinder and support brackets were made out of carbon fibers. Dummy aluminum substrates were then mounted (see Fig. 7.9) to the cylinder using the brackets on a coordinate measuring machine table. Known loads were then applied to the substrate and the deflection of the brackets were measured. Good agreement with the FEA results were obtained.

Work has also started on the vacuum system design. One of the first things that needs to be done is to understand the gas load. A model comprised of about 5% of the BTeV Pixel Detector (in terms of surface area) was built for the purpose of measuring its gas load due to outgassing and to understand how the gas load affected the ultimate vacuum pressure of the chamber. The model consisted of six substrates with dummy modules. A carbon-fiber shell supported the substrates. Kapton strips simulated the electrical flex cables. An aluminum plate served as a cable strain relief plate and a heat sink. The test was set up so that the model and the cable strain relief plate/heat sink was each cooled independently. Fig. 7.10 shows the model.

When the model and heat sink were at room temperature, the vacuum pressure was 3.4×10^{-7} torr and the gas load was 5.2×10^{-4} torr-L/sec. Cooling the model and heat sink to -10°C cut the gas load and the vacuum pressure in half. By cooling the heat sink to -160°C , the vacuum pressure was brought down to 1.0^{-8} torr. Analyzing the residual gas analyzer (RGA) readings at each temperature, it was found that water vapor was the main load and that cooling the heat sink to -160°C resulted in the heat sink acting as a cryo-panel that pumped water at a rate of 19,000 L/sec. Thus, using the cryo-panels in conjunction with other pumps such as cryopumps to pump on non-condensable gases can result in the pixel vacuum vessel's ultimate pressure to be $< 10^{-7}$ torr, which is the minimum acceptable pressure in the beam regions.

Prototype printed circuit boards (the feedthrough boards) which carry the signals from

inside to outside the vacuum vessel have also been tested and the results validate our conceptual design. Progress has also been made in testing various ideas of the rf shield. These include the use of aluminum foils, Be/Cu wires, and stainless steel strips.

To check the robustness of the high density flex circuits after multiple flexes due to the movement of the pixel detector in and out the beam, cable flexing tests have been started. Initial results look promising. To address question of the cables passing thorough a cryopanel, the effects of cold temperature on the electronic flex cables have ben tested. Prototype signal and power cables have been repeatedly flexed while immersed in liquid nitrogen and with current running through them (10 mA for signal cable, 1.5 A for power cable). The results also look promising.

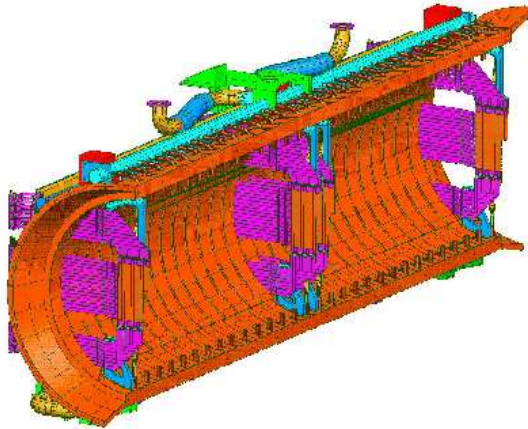


Figure 7.8: Side view of the vacuum vessel and support structure for the pixel detector.

7.1.7 Future R&D Plan

In the 1999 beam test, we demonstrated the construction and operation of single-chip pixel detectors. The results also showed that the resolution achieved with such a detector is better than $9\text{ }\mu\text{m}$ at all angles of incidence within the acceptance of the BTeV spectrometer. We plan to carry out another series of beam tests over the next two years starting this summer.

The primary goals of the beam test are to measure the efficiency, charge sharing, and resolution as a function of the irradiation previously accumulated on pixel detectors, under various conditions of incident angle of the track, threshold setting, and bias voltage of the sensors and with readout chips from various vendors.

We plan to build, install, and operate a pixel detector telescope based on 6 to 8 single-chip pixel detector planes. This will provide us with useful operational experience of a small scale pixel detector system. We also plan to test a 5-chip pixel module in the test beam.

Prototype mechanical and cooling support structures made out of TPG have been ordered and their mechanical and thermal properties will be evaluated.



Figure 7.9: Pi
carbon fiber o

s made out of



Figure 7.10: 5% model of the BTeV pixel detector, with dummy silicon modules assembled on six Al substrates.

Our next step is to address system issues. From our test beam experience, as well as operational experience from other large experiments, systems issues such as power supply, grounding, cabling, and connectors are potentially the most problematic areas. To understand and address these system issues, we plan to assemble a 10% system. With such a system, we can also carry out a thorough investigation of a complete electrical, mechanical and cooling system.

This will also enable us to operate a small system in the real C0 environment in a prototype Roman pot structure. Issues such as rf pickup, shielding, and fast readout coupled to a prototype trigger processor can be studied in detail there. At the same time, it will

allow us to understand a lot better the yield at the various steps of production, as well as how to assemble reliably the full scale pixel system.

Bibliography

- [1] J.A. Appel et al., Performance of prototype BTeV silicon pixel detectors in a high energy pion beam, FERMILAB-Pub-01/229E, to be published in *Nucl. Instrum. Meth. A*.
- [2] F. Lemeilleur et al., 3rd RD48 Status Report, CERN/LHC 2000-009 (31 December 1999).
- [3] ATLAS Pixel Detector Technical Design Report, CERN/LHCC 98-13.
- [4] M.R. Coluccia et al., Characterization of prototype BTeV silicon pixel sensors before and after irradiation, FERMILAB-Conf-01/344E.
- [5] D.C. Christian et al., Development of a pixel readout chip for BTeV, *Nucl. Instru. Meth. A*435 (1999), 144.
- [6] P. Jarron et al., 3rd RD49 Status Report Study of the Radiation Tolerance of ICs for LHC, CERN/LHCC 2000-03 (13 January 2000), D.C. Christian et al., FPIX2: A Radiation-hard Pixel Readout chip for BTeV, FERMILAB-Conf-00/316E.
- [7] J. Hoff et al., PreFPIX2: Core Architecture and Results, FERMILAB-Conf-00/260E.
- [8] G. Chiodini et al., Radiation tolerance studies of BTeV pixel readout chip prototypes, FERMILAB-Conf-01/214E.
- [9] G. Chiodini et al., Single event effects in the pixel readout chip for BTeV, FERMILAB-Conf-01/369E.
- [10] S. Cihangir et al., Characterization of indium and solder bump bonding for pixel detectors, *Nucl. Instru. Meth. A*476 (2002), 670.
- [11] S. Kwan et al., A Study of Thermal Cycling and Radiation Effects on Indium and Solder Bump Bonds, FERMILAB-Conf-01/377E.
- [12] S. Zimmermann et al., Development of high data readout rate pixel module and detector hybridization at Fermilab, *Nucl. Instru. Meth. A*465 (2001), 224.
- [13] S. Zimmermann et al., Development of a High Density Pixel Multichip Module at Fermilab, FERMILAB-Conf-01/247E.

7.2 Silicon Forward Tracker R&D

The R&D program for the Silicon forward tracker is organized along four main directions:

- development/choice of the Silicon microstrip sensors;
- development of the front-end read-out electronics;
- development of the detector support mechanics and the embedded cooling system;
- development of the Data Acquisition System (DAQ) for use in bench tests, beam tests, and quality assurance.

We report here on progress achieved thus far in each of these R&D directions.

7.2.1 Silicon microstrip sensors

Two groups are working in this area, namely INFN-Milano and the University of Tennessee. The first goal of this R&D effort is to find out if one of the microstrip sensors currently available on the market, is suitable for BTeV applications. We are presently considering the sensors developed by Hamamatsu and ST Microelectronics (France) for the CMS experiment. They are very promising and simple since they are of the p/n type and feature a very high breakdown voltage, typically around 800 V, which can ensure a full depletion even after radiation-induced type inversion. They should meet the requirements for BTeV.

We are characterizing samples in Milano (see Fig 7.11) and in Tennessee (see Fig 7.12) using two identical test stations developed by the IDE-AS company (Oslo, Norway) for this kind of application

The stations, which include a fast-preamplifier and a read-out section, have several features, which allow the user to change the preamplifier electrical parameters to permit a full characterization of the sensors. We have successfully installed the two stations and are now testing the sensors with a β -source and an infrared laser. The two sensors we are working on in Milano produced by ST, are $9 \times 9 \text{ cm}^2$ wide and are assembled in a daisy-chain configuration (see Fig 7.13 and Fig 7.14), identical to that foreseen for BTeV. This allows us to study the performance of long strips. The sensors under investigation in Tennessee are single sided versions of Hamamatsu S6934. These devices have 640 strips with a $50 \mu\text{m}$ pitch on 3.2 cm square wafer. The read-out strips are AC coupled. In Tennessee, we are preparing to test the wider sensors produced by Hamamatsu for CMS.

The infrared laser source we assembled in Milano can be focused into a few micron wide spot by means of a lenses and can be moved by actuators along X and Y -directions (see Fig 7.15). This system will be crucial in speeding up all the quality checks on the sensors during the final production phase.

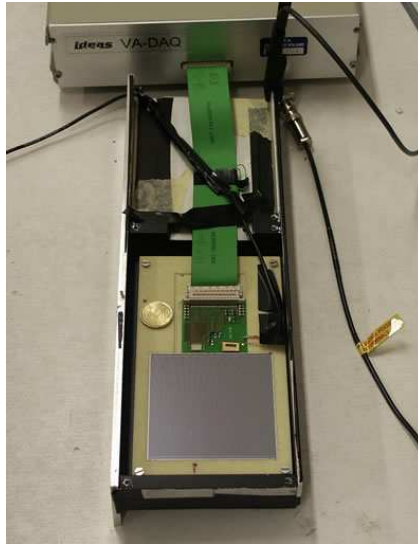


Figure 7.11: A prototype microstrip detector attached to the IDE-AS system box in Milano

7.2.2 Front-end read-out electronics

The aim of the project is to develop a new read-out chip which has to be very fast, low noise and data driven. The chip has to sustain rather high radiation doses and to read out long strips having a typical capacitances of the order of 20 pF. For this chip we plan to use a digital section based on an architecture similar to what we developed for the pixel read-out chip.

Two steps have been already accomplished in this effort: the choice of the technology and the definition of the basis circuitry for the analog channel. The chosen technology for integration is a deep sub-micron process, which can be made highly radiation resistant with some proper layout prescriptions. In particular we are considering the TSMC (Taiwan) process with 0.25 μm minimum feature size, which has been successfully used by other groups.

The basic scheme of the analog channel, together with some particularly significant test structures, has been completely simulated and is shown in Fig 7.16. The channel includes a charge-sensitive preamplifier, a second-order semi-Gaussian shaper and a comparator to provide 1-bit information, generating a logical 1 at the output if a signal exceeding the comparator threshold is detected. Fig 7.17 shows the waveforms (from Eldo simulations) at the output of the three circuit blocks as a response to an input current signals of short duration, corresponding to a 4 fC charge delivered by the detector. The corresponding files are now available and ready for the preparation of the layout. We have evaluated the costs for a preliminary submission and have already received offers for the CAD work for the layout.



Figure 7.12: A panoramic view of the sensor test station at the Tennessee University

We are now ready for the submission of the analog channel prototype and the related test structures.

The Milano and Pavia groups are working on this project with P. F. Manfredi as general coordinator.

7.2.3 Detector support mechanics and cooling

This R&D project is logically divided in three parts: the development of the micro-strip inner support, which provides a local support for the microstrip planes and all their components; the development of the micro-strip outer support, which holds the inner support; and the development of the cooling system for the front-end read-out chips. The first part is carried on by the Milano group, the second one by the Frascati and Milano groups, and the third one by the Bergamo group.

The main goal of this effort is to design a very light structure in order to minimize the amount of material inside the tracking volume.

We have already developed a prototype design for the inner support in collaboration with PLYFORM, a Milano based company with long experience in this sector having built parts of the support structures for the CDF and ATLAS vertex detectors. The support structure (see Fig 7.18 and Fig 7.19) is realized with carbon fibers and allows for a very precise relative positioning of all the components belonging to a same station. PLYFORM will produce a mock-up structure for tests of functionality shortly. By the end of the year we will have the complete prototype support structure in carbon fiber for a whole station.

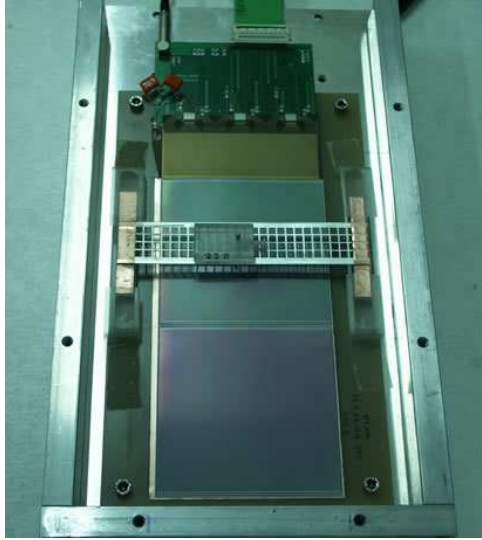


Figure 7.13: The two ST sensors assembled in a daisy-chain configuration. The sensor pitch is $183\ \mu\text{m}$ and the total number of strips is 512. At the bottom of the figure the read-out chips and the fan-out circuit are visible.

At this time we consider as a base-line solution for cooling a water/glycol mixture circulating in a duct embedded into the support structure, but we are also investigating other interesting possibilities. The Bergamo group has started the simulation work for several different cooling systems. They are investigating the performance of each system as well as the deformations induced in the mechanical structure. It is our intention, once the simulation work is finished, to assemble prototype systems to validate the simulations and to test their cooling performance.

The R&D activity on the outer support has just started. Several ideas have already been discussed and one in particular seems very promising. Instead of using a separate support for strips, we are considering the possibility to integrate it directly into the nearby straw structure. The idea is to reinforce, by means of carbon fibers, the straw structure in the region covered by the strip inner support in such a way that it can be made rigid enough to hold an entire strip station. The straw tubes could also benefit from this additional rigidity, since their own support structure could, in principle, be simplified and made lighter in the central region. Even the strip signal and power cables could be integrated in the same structure by simply depositing on it the required amount of copper stripes. We are working on this and plan to prepare a complete design in the near future.

7.2.4 DAQ system

The R&D activity for the design of a DAQ system, which can be used for tests of the Silicon Strip Detectors both in bench tests in the lab and in beam tests, has already reached



Figure 7.14: A magnification of the previous figure to show details of the read-out chips and the fan-out circuit.

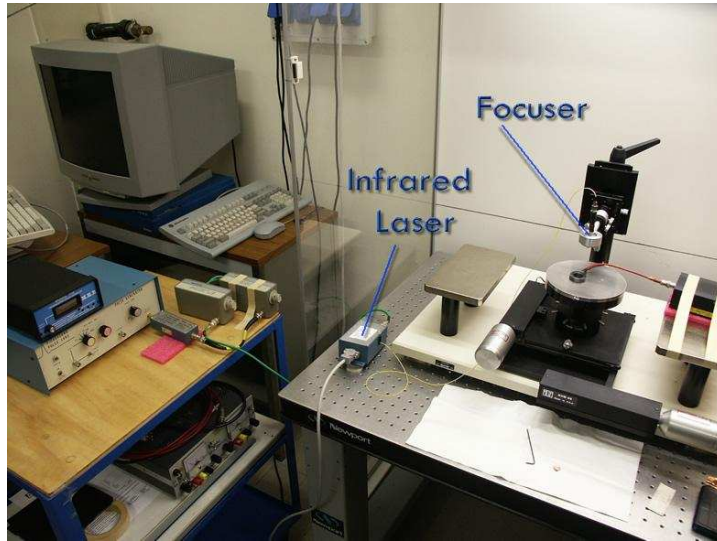


Figure 7.15: The infrared laser station with the X/Y moving chuck.

significant milestones. We are sure that this experience can be profitably used to design the future DAQ for BTeV.

The DAQ design is based on the PCI bus protocol, a widely used standard in the computing industry, which offers several benefits, one being relatively low cost. The digital part of the Silicon Strip Detector front-end is designed to be practically the same as that of the Pixel Detector, thereby allowing for a common read-out scheme for these two detectors. In our design each detector is connected to a PCI board containing an FPGA (for logic control)

**Analog channel
for the BTeV Forward Silicon Tracker**

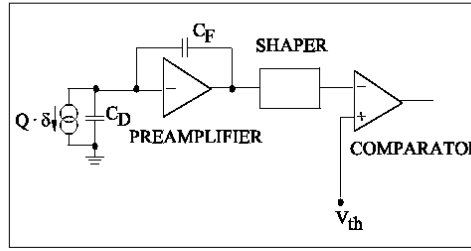


Figure 7.16: Basic scheme of the analog channel for BTeV microstrip front-end

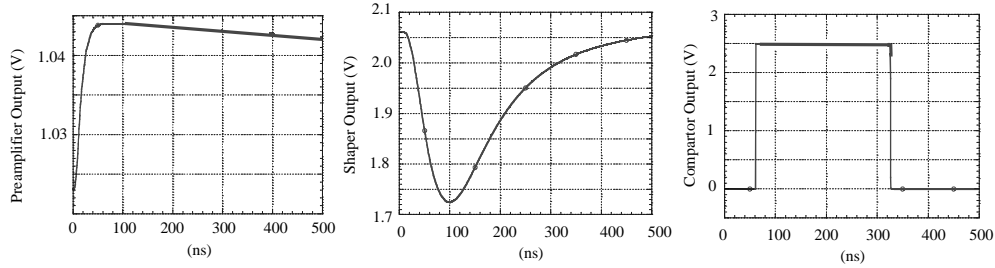


Figure 7.17: Waveforms at the output of the three circuit blocks of Fig 7.16 for 1 Minimum Ionizing Particle as simulated by ELDO.

and two 1 Mb memories; several PCI boards are lodged together on a PCI bus extender and finally connected to a host DAQ PC. Each time a strip (or a pixel) has data above threshold, the address along with pulse height and time-stamp information are sent to the PCI board to be stored on one of the two local memories. The FPGA's are programmed to handle the swapping between these two local memories and synchronization with the external read-out process (running on the host DAQ PC) in such a way to smoothly handle a sustained data rate, adequate to the beam test requirements. The central idea in this design, is keeping the event-builder algorithm as simple as possible, since an event, defined by all hits marked by equal time-stamp, is spread out over several PCI boards which can in principle receive data at different rates. In absence of a specifically defined strategy to synchronize the flushing of these memories, this sparse read-out scheme generates events spread-out over large chunks of memory, making the event builder extremely cumbersome. We have therefore designed an elegant mechanism to restrict the components of an event to be contained in a limited amount of memory, taking advantage of our ability to program the

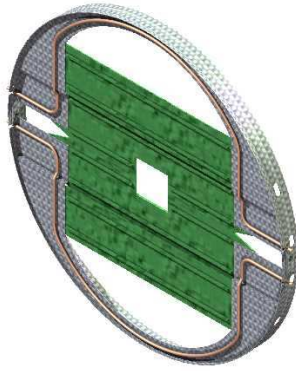


Figure 7.18: 3D sketch of the inner support for a single silicon plane. The back side of the plane is shown together with the cooling ducts. In grey is the carbon fiber structure, in green the ladders of silicon sensors and in red the cooling ducts.

FPGA to generate interrupt signals. The principle of operation of this read-out scheme is the following:

- data are received from a detector by the corresponding PCI board in one of its two internal memories
- as soon as any memory in the system is full, all boards are synchronously commanded to swap their memories. The ones used so far are frozen and immediately read-out to the host computer, while the others are used to continue reading events from the detectors without any data loss.
- events are fed to the host computer on a statically allocated shared memory, acting as a compensating buffer to allow for unexpected data rate fluctuations.
- data are then continuously flushed from this memory to disk by a consumer process, building events on the fly.

This is an event-driven scheme: data are collected as soon as they are produced by a detector, and no burden is placed on the DAQ software to generate signals to start a read-out chain. This is important, since it allows testing the full functionality of the detector in an environment similar to the one envisaged for the final data taking, where no trigger is used to read-out events.

Several components of this read-out have already been implemented, on a Linux platform:

- the PCI board and the microprogramming of the FPGA to send and generate control signals and interrupts.

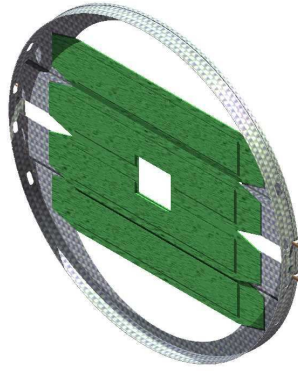


Figure 7.19: 3D sketch of the inner support for a single silicon plane. The front side of the plane is shown.

- an abstract interface to the underlying driver (currently we are using a licensed driver, produced by Jungo). This will allow us to switch to other drivers should competitive ones emerge in the future.
- the interrupt-handler processes, in charge of starting the read-out of a PCI memory, synchronizing the read-out of all other boards and transferring data to an external shared memory
- the read-out process, owner of the shared memory and responsible for synchronizing with the consumer process to flush events on a storage media
- the event-builder (a prototype)
- a package for message transmission among cooperating process (based on the native Linux IPC protocol)

We have already in place a code management system, based on CVS, that allows us to keep track of successive releases of the code, making it possible to allow concurrent development at all participating institutions.

Future developments of this project encompass several key aspects:

- a sophisticated error detection system, taking into account both faults generated by the hardware (such as broken components or loss of synchronization) as well as those generated by beam conditions (e.g., memories swamped by surge in beam intensity); and
- a set of GUIs to allow users to drive the DAQ system.

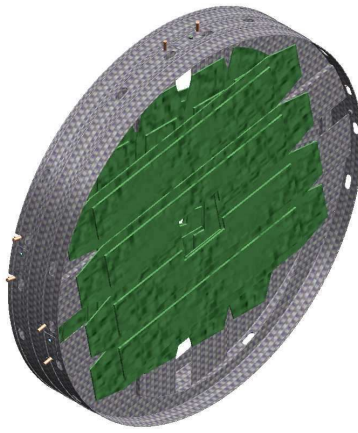


Figure 7.20: The support structure is such that the three views can be stacked to form a station. Reference pins guarantee a very precise relative alignment when assembling the components of the structure.

7.3 Straw Tube Tracker R&D

The Straw Tube Tracker is part of the forward tracking system. The major functions of the forward tracking system are to provide high precision momentum measurements for tracks found in the pixel system, to reconstruct tracks that do not pass through the vertex detector and to project tracks into the RICH and EM Calorimeter. The baseline system consists of 7 stations (1 arm). Most of the solid angle is instrumented using straw tubes. The high occupancy region near the beam is instrumented with silicon microstrip detectors. Each Straw Tube station consists of 3 views and each view has 3 close-packed layers. The following is a summary of our recent R& D work; details can be found on our website [1].

7.3.1 Single Straw Prototype

In order to become familiar with the operation of a straw tube detector, we obtained an old straw chamber from the Focus experiment and set up the hardware and software necessary to read out cosmic ray data. This included a flammable gas system so we could test various gas mixtures, high voltage supplies, low voltage supplies, CAMAC based TDCs, scintillation counters for triggering and a LabView based DAQ.

We have about 1000 straws that were designed at the University of Indiana for BTeV, based on the design of the Atlas straws. These are made of two layers of kapton. The inner conductive layer is carbon loaded kapton with a thin film of aluminum deposited on it. The outer layer is plain kapton.

We constructed several single straw prototypes using these straws. This required designing end plugs to attach the straws to an end plate. The anode wires are then attached to an insulated anode board, inside a gas volume.(see Fig 7.21)

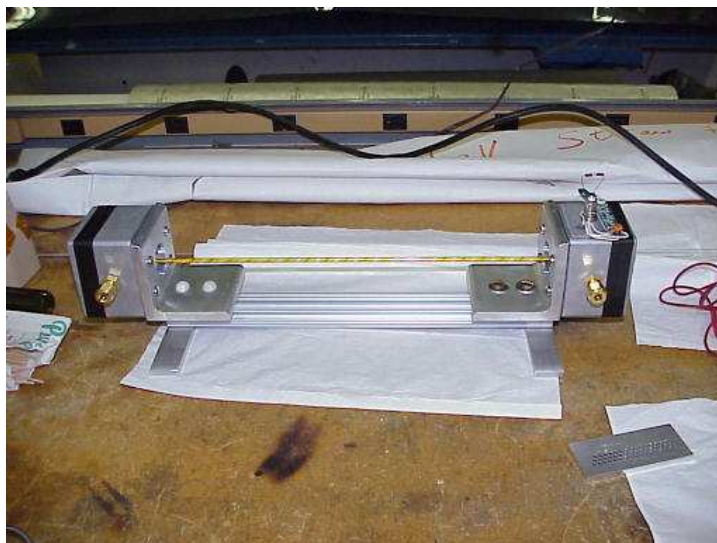


Figure 7.21: Single Straw Prototype

The anode wire is kept centered in the straw with helical structures called “twisters”. For short straws one twister at each end is sufficient but straws longer than 80 cm need additional twisters. We obtained twisters from 3 different manufacturers. It is important that the twisters have smooth surfaces in order to not snag the wires. The twisters obtained from A.F. Leis in Columbus, Ohio were of good quality. Those from the other manufacturers needed extra smoothing. We used the electronics from the FOCUS chamber to read out these prototypes.

Gas gain measurements were done with an ^{55}Fe source and a 50:50 Ar/Ethane gas mixture. The University of Houston group has started doing ageing tests using a 100 mCi ^{90}Sr source. These studies will influence the final choice of gas and also test for problems that might occur with outgassing from materials such as the glue that we use to keep the twisters fixed in the straws.

7.3.2 Wire Stringing and Tension Measurement

The 25 μm diameter, gold-plated tungsten anode wires are inserted into the straws by blowing a lead wire through, tying the anode wire to the lead wire and pulling the lead wire back. A weight is hung from the wire to apply the correct tension. We have successfully strung wires through straws up to 4m long without any problems. The setup is shown in Fig 7.22.

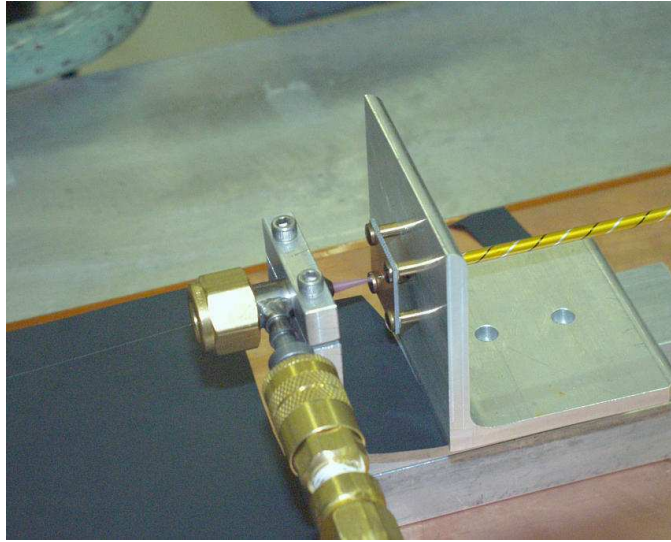


Figure 7.22: Nozzle for wire stringing

The tension on the anode wire needs to be approximately 50g. During production it will be necessary to have an efficient method to check the tension of each wire as it is strung. With the aid of an undergraduate student from UC Davis (Long Pham) we have set up an acoustic excitation method for measuring tension – a loudspeaker is used to induce a mechanical resonance in the wire. A potential of 50V is put on the wire so the capacitance variation induces a signal. The current setup uses a LabView data acquisition system that

generates a variable frequency to drive the loudspeaker and plots the response of the signal as a function of frequency. The resonance is easily observed. We are working on developing a faster system.

In order to reduce the occupancy, the straws are read out at both ends, the anode wires being separated in the middle by a fused glass capillary tube. The group at the University of Virginia is planning to build a wire fusing machine based on the design of the one used for ATLAS.

7.3.3 Straw Stretch/Tension Tests

In the summer of 2001 we had a high school teacher, Jan Dudzik, set up a system to test the response of the straws to changes in temperature and humidity. Three straws approximately 1m in length were hung vertically with a weight attached and placed in a box where we could control the temperature and humidity. The change in length of the straws was measured with inductive proximity sensors. The results are shown in Fig 7.23. The straw responded very slowly to changes in humidity, probably because the straw was sealed which will not be the case when running the experiment. We decided that it would be more realistic to fix the length of the straws and measure the change in tension as a function of environmental conditions.

This has been done using a load cell to measure the tension. The results are shown in Fig 7.24. A straw was set up in a frame and stretched until it had a tension of about 190g. The tension rapidly decreased to about 160g. The relative humidity was then decreased from 37% to 0.5% and the tension increased to 260g. When the humidity was increased again to 37% the tension dropped to 120g. We cycled the humidity several times between 37% and 0% and the tension varied from 120g to 260g as expected, showing that the initial reduction in humidity stretched the straw irreversibly but was stable after that. From these tests we have learned that it will be necessary to run the straws under well controlled environmental conditions. We plan to flow dry nitrogen in a volume surrounding the straws to keep the humidity close to 0%. It is necessary to apply a tension of about 200g to each straw in order to keep it straight. We have the straw now held under zero humidity and are looking at long term changes in tension.

7.3.4 Two-Module Prototype Construction

The final BTeV Straw system will be built in “half-views”. Each half-view will have several modules where a module consists of 3 layers of 16 straws. We have constructed a 2-module prototype to be used in a beam test later this year. The construction of a module proceeds as follows: the straws are cut to the correct length and twistors are glued in; the 3 layers of straws are set up on a corrugated base to form a close-packed array; the end-plugs are inserted into the end-plates and silver epoxy is injected into the region around the plugs; the end-plates are attached to the straws; finally the outside of the endplates is potted with structural epoxy. (see Fig 7.25)

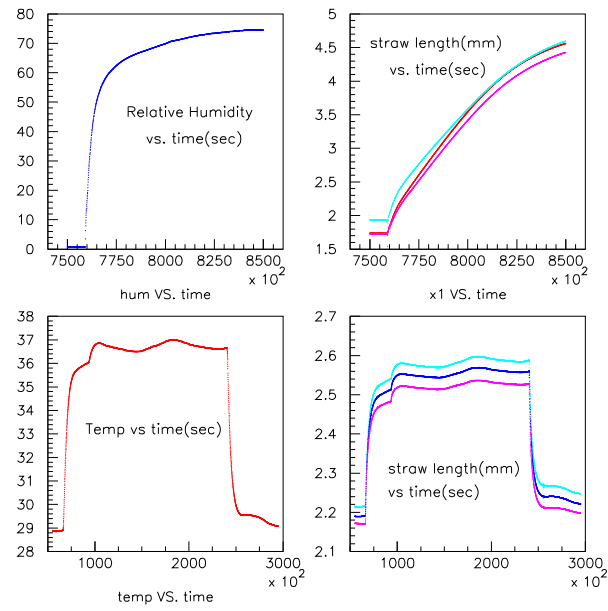


Figure 7.23: Straw Length dependence on humidity and temperature

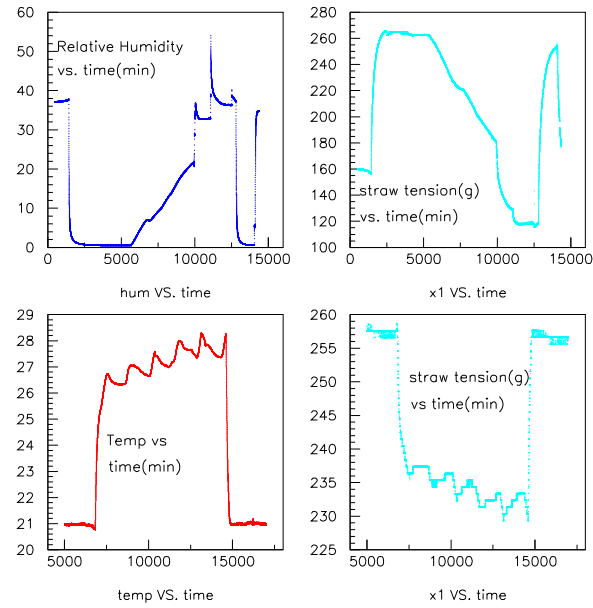


Figure 7.24: Straw Tension dependence on humidity and temperature

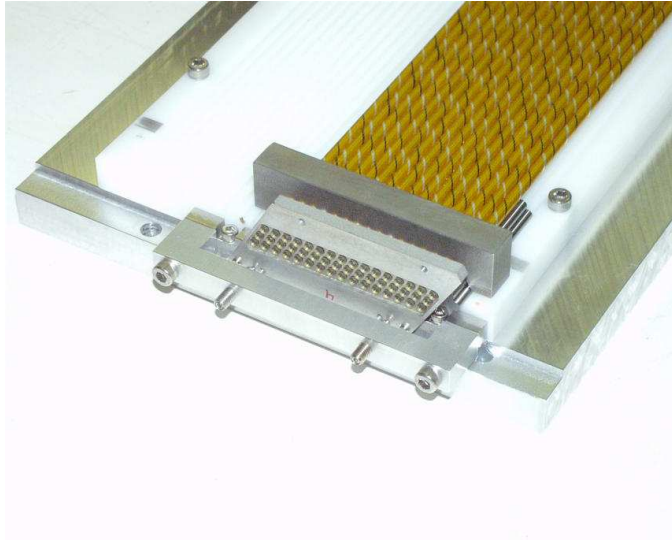


Figure 7.25: Straw module assembly

Anode boards were designed with connectors to COT cards (from the CDF Central Open-cell Tracker). Each COT card has three 8-channel ASDQ chips ([2]). The output of the COT cards is converted from LVDS to ECL via a translator board so we can use existing LeCroy TDCs. We have discovered several problems while constructing these modules: (1) The straws are not perfectly round which makes it difficult to have them close-packed without bowing outwards. (2) About 20% of the straws have leaks. This was unexpected as we had tested a few straws early on and had not found any leaks. The leaks are associated with places where there appears to be a lack of glue between the two layers of kapton. (3) Silver epoxy was used to attach the anode boards to the module end-plates and shorted out a couple of traces on the anode plate.

We believe the problems of out-of-roundness and lack of sufficient glue are solvable as we have done measurements of the straws to be used for the CKM experiment and found that they do not have these problems. We plan to order more straws with stricter quality control on the construction.

7.3.5 Other R&D

If an anode wire breaks it will draw a lot of current and probably trip the HV power supply. As each HV channel will supply many straws we are investigating the possibility of using a fuse for smaller sets of straws. The group at Southern Methodist University (SMU) has identified a commercially available fuse that is possibly suitable for this application. They have obtained some samples and are conducting tests. We have had discussions with Oak Ridge National Laboratory about the design of the front-end boards and they have some seed money to do preliminary design work.

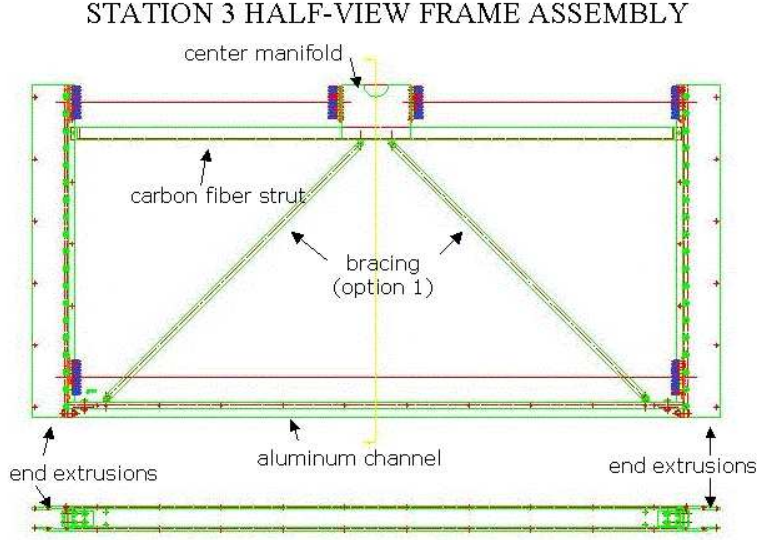


Figure 7.26: Half-View Frame

We have made a detailed design of the flammable gas system to supply the BTeV straws [3]. The group at SMU is designing a gas gain monitoring system.

We are planning to do radiation damage tests on all materials that will be used in the straw tracker.

We have a preliminary design of the half-view frame that will be used to support the straw modules, shown in Fig 7.26. A finite element analysis has been done to examine the flexing of the frame due to the tension on the straws. A mechanical model of the frame for a station 3 chamber is under construction. This will be put under the tension in order to test that the calculations are realistic.

The group at Frascati is developing a method of accurately tracking the alignment of the straws with respect to the silicon microstrips.

7.3.6 Simulations

We have continued simulations of occupancy and radiation levels with the updated BTeV spectrometer. Drift time and gas gain calculations have been done using Garfield. We are refining tracking code to reconstruct tracks in the forward tracking system in order to check the effect of any design modifications.

Bibliography

- [1] “BTeV Straws Web page”, <http://www-btev.fnal.gov/atwork/detector/straw>
- [2] “The ASDQ ASIC for the Front End Electronics of the COT”, W. Bokhari et al., <http://penn01.fnal.gov/cot/doc/ASDQ-new.ps>
- [3] “BTeV Straw Gas System”, Terry Tope, <http://www-btev.fnal.gov/cgi-bin/DocDB/ShowDocument?docid=259>

7.4 Ring Imaging Cherenkov Detector R&D

The RICH detector consists of a gas radiator (C_4F_{10}) followed by a spherical mirror focusing Cherenkov light onto arrays of Hybrid Photo-Diodes (HPDs). In the initial design, we planned to have an aerogel radiator in front of the gas radiator to improve low momentum particle identification. Cherenkov photons from both radiators were detected on the same array of HPD sensors. Based on detailed simulation studies that showed the small number (~ 10) of aerogel photons would be obscured by the overlapping rings of photons from the gas radiator, we have decided to replace the aerogel radiator with a liquid radiator (C_5F_{12}) that has its own Cherenkov light detection system consisting of traditional photomultiplier tubes (PMTs).

7.4.1 HPD Development

In an HPD, a photo-electron emitted by the photo-cathode is accelerated onto a segmented silicon diode by a very high voltage (~ 20 kV). The BTeV HPD is based on the cross-focusing tube developed by DEP in the Netherlands (Delft Electronic Products B.V.) working with the late Tom Ypsilantis and Jacques Seguinot as part of the LHCb group. Initially these tubes were developed with 61 silicon pixels inside. Much finer segmentation is needed in the LHCb RICH detectors, thus they plan to have 2048 pixels in each HPD tube. This makes it impossible to use the original DEP design in which pixel signals are individually transferred outside of the tube by means of pins. Therefore, the LHC-b group is developing their own silicon chip in which the diode is integrated with the front-end readout and both reside inside the tube. The diode segmentation needed in BTeV is 163 pixels per tube which represents a small modification of the original DEP scheme. This approach decouples development of read-out electronics from the HPD manufacture.

We developed the 163-channel HPD together with the DEP. The redesigned diode-package is shown in Fig. 7.27. Two tubes of this type have been manufactured by DEP and successfully tested at Syracuse. The pulse height spectrum for one of the HPD channels, obtained with low intensity LED light and VA-RICH readout electronics (adopted from the CLEO-III RICH) is shown in Fig. 7.28. Peaks due to one, two, and three photo-electrons reaching the same pixel within the integration time are observed. We have ordered 13 more HPDs, to be delivered by early Fall 2002. They will be used in a beam test of a prototype of the RICH later this year.

7.4.2 HPD Front End Readout

The HPDs will be equipped with binary readout. Each channel will have a discriminator with its threshold individually adjusted with an internal DAC. A Cherenkov photon detected by the HPD produces on average 5000 electrons. To achieve a signal to noise ratio of about 7:1, we are aiming at front-end electronics with equivalent noise charge of 700 electrons. We are working with IDE AS from Norway to develop such electronics. The Syracuse group

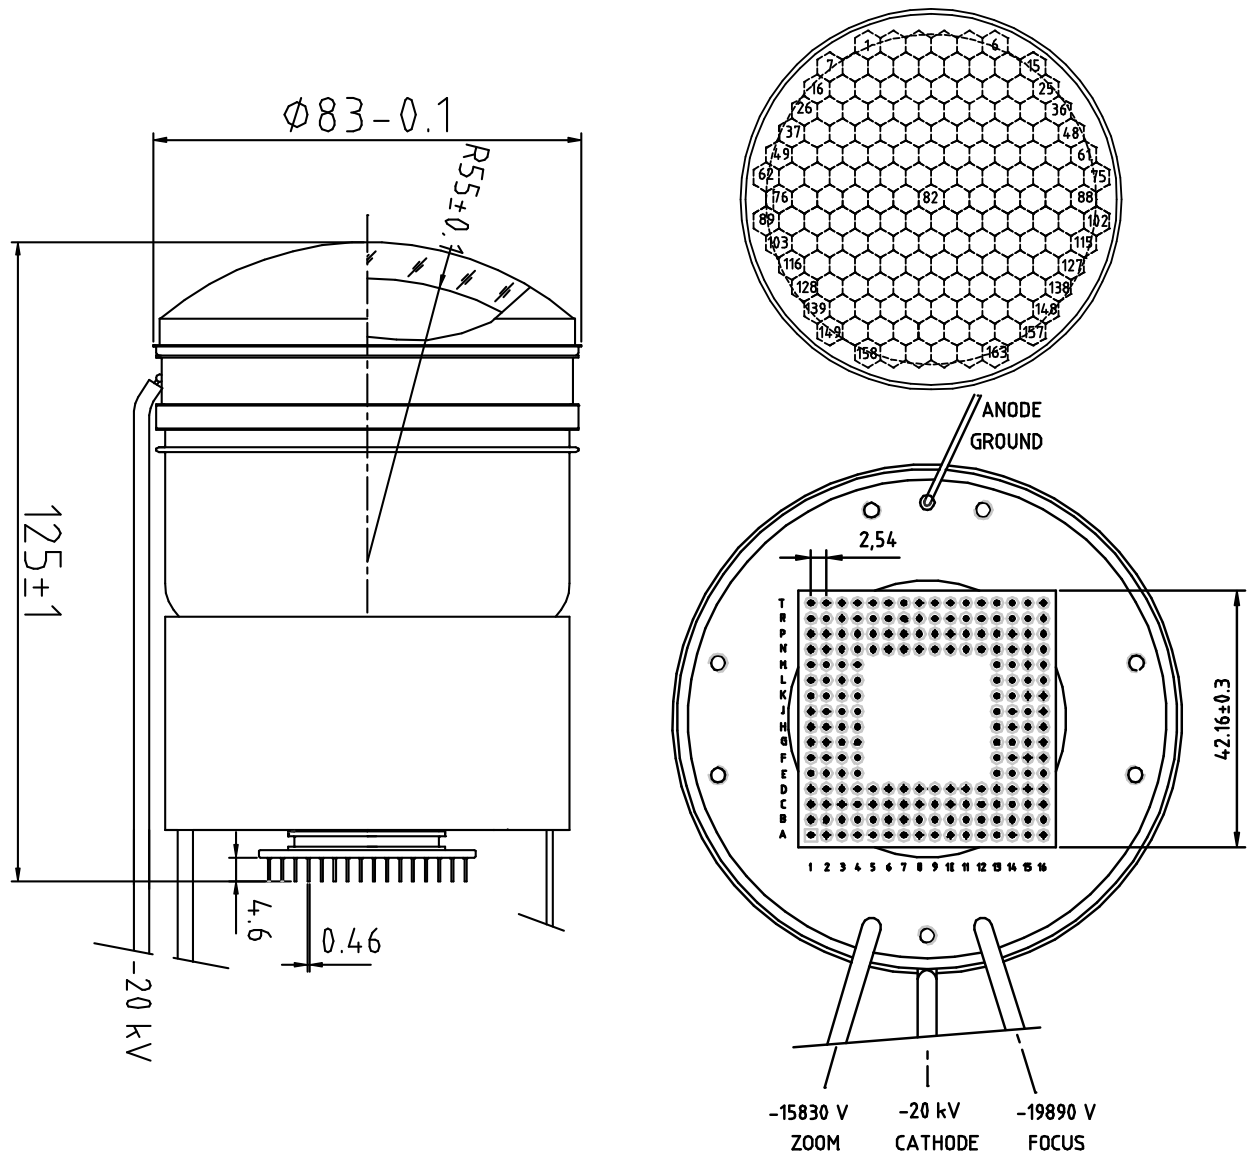


Figure 7.27: BTeV HPD (DEP PP0380AT). The dimensions are given in mm. The 163 silicon pixels on the cathode are shown in the upper right drawing. The layout of pins penetrating the tube is illustrated in the lower right picture.

previously worked with this company on development of a custom-made ASIC called VA-RICH and of associated front-end hybrid boards that were used in reading out the CLEO-III RICH detector. A different adaptation of the VA chip family has been produced for the BTeV HPD. We refer to this new ASIC as VA-BTEV. The VA-BTEV chip amplifies, shapes and discriminates the signal. It has a fast peaking time (72 ns) matched to Tevatron bunch crossing time (132 ns). The fall time (200 ns) extends the processed signal to the next bunch crossing, with negligible loss of Cherenkov light thanks to the small channel occupancy. Since each chip has 64 channels, a front-end board will house three of them. They will be connected to the HPD output pins via a small interface board. This analog part of the front-end hybrid

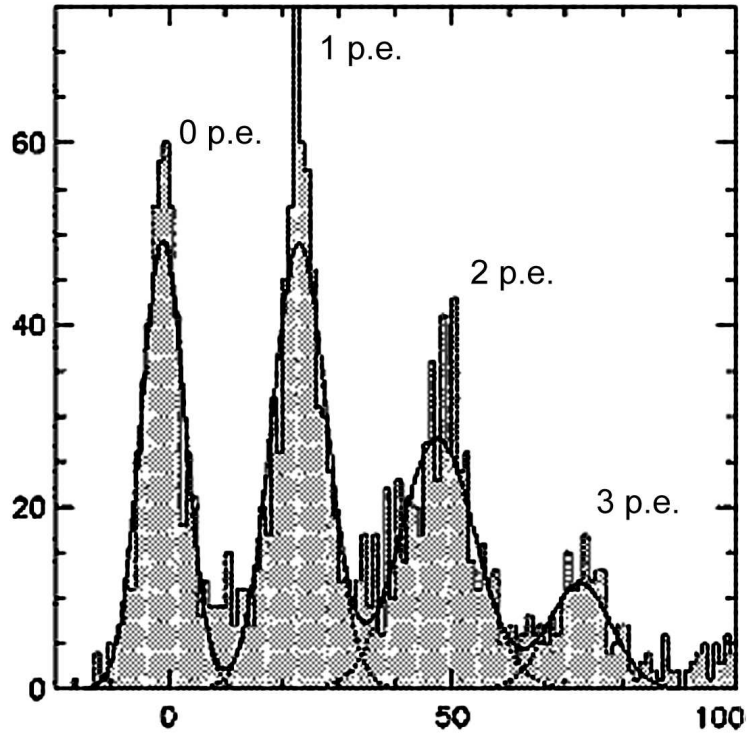


Figure 7.28: Pulse height spectrum obtained for BTeV HPD with VA-RICH readout. The pedestal peak is positioned at zero. The subsequent peaks corresponds to one, two and three detected photo-electrons.

will be well isolated from the digital part for the best signal to noise performance. Binary signals for each channels are fed in parallel into the digital part hosting an FPGA, which serializes the output, encodes channel address and attaches a time stamp.

The layout of the prototype VA-BTeV front-end boards is shown in Fig. 7.29. The first prototypes are now being tested on a bench using the DAQ system that later will be used in the RICH beam test. The next iteration of the design will have a flex circuit between the analog part (mounted directly to the HPD) and the digital part. This geometry is necessary to allow for close packing of HPDs. The next generation of the prototypes will be used in the beam test.

7.4.3 HPD High Voltage

The HPD requires application of three separate HV levels. The tube is sensitive to the stability of this applied power. We are using a relatively high noise power supply (Acopian) with 1 V of ripple, but RC filtering near the HPD renders this noise relatively harmless. We also have supplies with only 3 mv of ripple (Matsuada) that could, in principle, work without filtering, though we are concerned about line pickup. We have determined that ripple greater than 10 mV at the HPD does degrade the signal-to-noise performance. In a

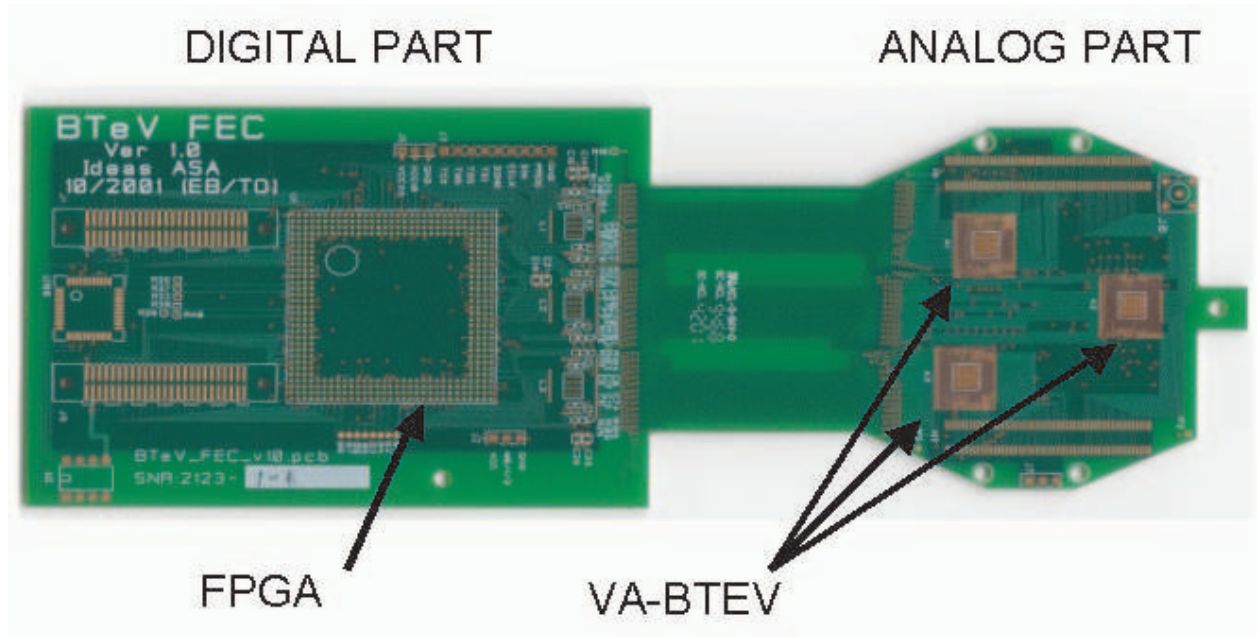


Figure 7.29: Layout of the first prototypes of BTeV RICH front-end hybrid board.

full-scale distribution system, we may need to employ a low ripple supply as well as local filtering. Additional tests are in progress.

7.4.4 Mirrors

Spherical mirrors at the end of gas radiator focus Cherenkov photons onto the HPD arrays. The mirror curvature must have a radius of about 700 cm to match the RICH dimensions. Individual mirror tiles will have a hexagonal shape with a diameter of 60-80 cm. Mirror quality must be good enough to focus 95% of a parallel beam of light within a spot of 3 mm diameter, in order to preserve the Cherenkov angular resolution. One relatively cheap and well established technology to produce such mirrors is to use a glass substrate. The glass mirrors would add about 5% of radiation length in front of the EM calorimeter and require a much heavier support structure. We are investigating ways to reduce amount of material in the mirrors by using some composite materials like carbon fiber or foams to partially or completely replace the glass. Mirrors as thin as 1% of a radiation length can be built. We are determining if they have acceptable quality at an affordable price.

The Torino group has ordered two mirror prototypes from Turnov CZ (COMPAS R&D Consortium); one made out of glass, and the other with a thinner glass substrate reinforced by carbon fiber. Their quality will be soon investigated on a test bench at CERN. Since these mirrors have the right size and the right radius of curvature, we plan to use them in the gas RICH test beam set-up. CMA in Tucson AZ produced a mirror prototype for us using a carbon fiber substrate replicated on a glass master. The first CMA prototype has

much shorter radius of curvature (190 cm), therefore it will be tested on a bench only. This technique was used to produce good quality mirrors for the HERMES RICH. The Torino group is collaborating with a local Italian company on developing a similar technique to reduce costs.

7.4.5 Liquid Radiator

The main limitation of the gaseous RICH is its inability to separate kaons from protons below the kaon radiation threshold (9 GeV/c) which dilutes same-side and away-side kaon tagging. The overall measure of flavor tagging, $\epsilon\mathcal{D}^2$, suffers by as much as 25% in the case of the B_s due to this effect. To fix this problem we previously proposed adding an aerogel radiator in front of the gas radiator, with aerogel Cherenkov photons detected in the same HPD array as used for gas photons. Such a dual radiator RICH has been in operation for last few years in the HERMES experiment. Initial simulations of the aerogel radiator in BTeV RICH showed that it did separate kaons and protons below 9 GeV/c at some level. These initial studies neglected backgrounds from minimum bias events and from photon conversions in the BTeV detector components (beam pipe, tracking system, RICH radiators). More realistic simulations proved that no useful K/p separation could be achieved with aerogel detector, even after the Aluminum beam-pipe was replaced with Beryllium, and reconstruction algorithms were improved to suppress photo-conversion backgrounds. The aerogel fails because it produces large and faint Cherenkov rings that overlap many intense gas radiator rings as well as other aerogel rings. Pattern recognition of these faint images is just not possible in busy $p\bar{p}$ events. It should be noted, that track multiplicity in HERMES experiment is 1 or 2, compared to about 80 (on average) particles in the BTeV RICH. Thus, positive experience with aerogel radiator in HERMES does not carry over into the BTeV event environment.

We are now proposing to replace the aerogel radiator with a liquid radiator. The selected liquid radiator, C_5F_{12} , has much larger refraction index than aerogel ($n = 1.24$ vs. $n = 1.03$) producing more intense Cherenkov rings even from a thin layer of liquid. The C_5F_{12} radiator was successfully used in the other experiments (e.g. DELPHI, and SLD). Unlike an aerogel radiator, in which the Cherenkov photon yield is limited to a constant value by light scattering in the aerogel itself, the Cherenkov yield in liquid scales with radiator thickness. The large refractive index makes the Cherenkov rings even larger, which is actually beneficial, since optical paths of liquid photons now decouple from optical paths of gas photons. About two thirds of liquid photons hit the side-walls of the RICH where they can be directly detected without any focusing elements (so-called “proximity focus”). The fraction of gas photons striking the side-walls is essentially negligible as they are all incident at the mirror positioned at the end of the gas radiator. Thus, the two main limitations of aerogel radiator, faintness of the image and contamination from gaseous rings, are eliminated. At the same time, the refractive index of C_5F_{12} is low enough for kaon and proton Cherenkov rings to have sufficiently different radii up to 9 GeV/c to be distinguished with relatively large diameter PMTs instrumenting the side-walls. It should be also noted that although some

liquid photons do reach the mirror (about one third), they do not contaminate gas images, since they are imaged outside the instrumented part of the HPD planes.

Simulations of the liquid radiator performance for a sample of low momentum (< 9 GeV/c) kaons and protons are compared to the simulations of the aerogel radiator in Fig. 7.30. For aerogel (top picture) the distribution of protons in $\chi^2(K) - \chi^2(p)$ is essentially indistinguishable from the distribution obtained for kaons. (Here, each χ^2 is defined as $-2\log(\text{Likelihood})$ as described in chapter 4.) For the same sample of events and tracks, the liquid radiator (bottom

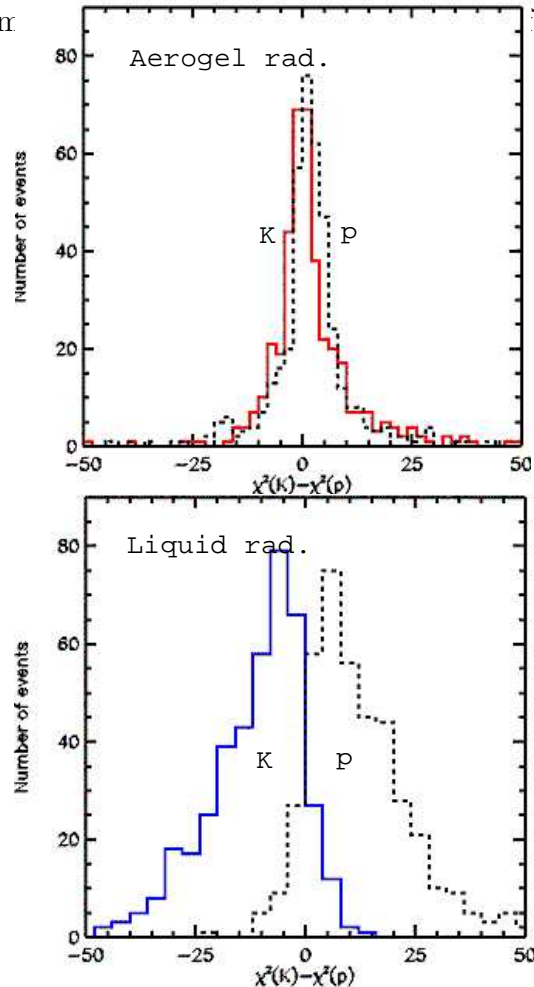


Figure 7.30: Performance of the RICH detector with aerogel (top) and liquid radiator (bottom) on a sample of low momentum (4-9 GeV/c) tracks. The underlying environment is a beam crossing which contained a $b\bar{b}$ event and on average two minimum bias events. Solid histograms show the kaon distribution, dashed histograms show the proton distributions.

We expect the chromatic error for C_5F_{12} radiator with photons detected in the visible range (wavelengths above 300 nm) to be 3.7 mrad per photon. This error should be combined in quadrature with a photon position error reflecting the size of the PMT face. Two-inch PMTs would result in a photon position error of 3.5 mrad and an overall single photon

resolution of 4.9 mrad. The number of PMTs and therefore the phototube cost can be cut in half by using 3" PMTs, with only 20% deterioration in Cherenkov angle resolution (5.3 mrad photon position error, 6.2 mrad overall). To cover the most illuminated parts of the RICH walls we will need roughly 5,000 3" PMTs per arm. With this coverage and with the liquid radiator 1 cm thick we expect to detect about 12.4 photons per track, resulting in a per track resolution of 1.88 mrad. Since at 9 GeV/c, kaon and proton Cherenkov angles differ by 5.34 mrad, separation would be 2.8 standard deviations for well-isolated Cherenkov images. Separation improves substantially for lower particle momenta.

Because of the large number of PMTs needed, minimizing cost per PMT is essential. The cheapest PMTs with single photo-electron capability are conventional head-on tubes, with 8-stage box dynode structure. With HV around 1 kV, their gain is of the order of a few times 10^5 and collection efficiency well above 90%. A standard bialkali photo-cathode with a borosilicate glass window provides a peak quantum efficiency around 30%. Dark count rate is orders of magnitude below the level that would impact RICH performance. At present, we are in contact with four different manufactureres making such photo-tubes in 3" size: Burle, Electron Tubes, Photonis and Hamamatsu. We are testing their PMTs in order to establish single photoelectron detection capability and efficiency loss in a magnetic field. We have already tested sample tubes from Hamamatsu and we are now testing Burle PMTs. We will soon test PMTs from Electron Tubes and Photonis. A single-photon response measured for R6233 PMT from Hamamatsu is shown in Fig. 7.31.

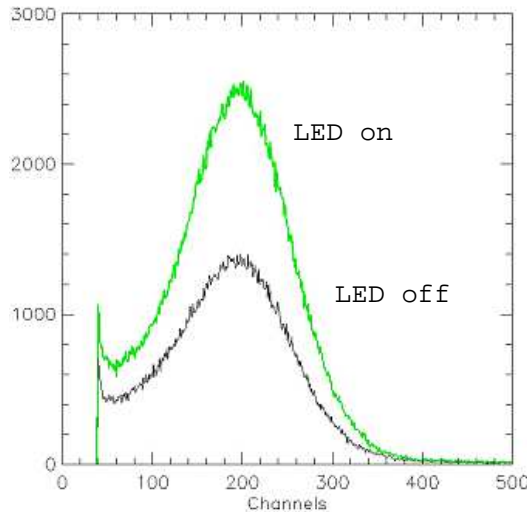


Figure 7.31: Pulse height spectrum obtained with R6233 PMT from Hamamatsu connected to MCA exposed to attenuated LED light (upper histogram). Dark counts (lower histogram) accumulated for the same readout live time. The single photo-electron peak is clearly visible with a good peak-to-valley ratio. Dark counts also show the single photo-electron peak because of thermionic emission of electrons from the photo-cathode.

It is possible that it will be cost effective to order PMTs with integrated voltage divider

boards, that could also serve for their mechanical support. Further R&D will be needed to identify a manufacture with the best price per performance ratio.

We have had initial discussions with IDE AS about front-end electronics which would have similar architecture as our HPD front-end boards but a simpler design. One front-end board would serve 64 PMTs, with signal cables soldered on both ends. Further work on PMT read-out scheme is needed.

We will also need to prototype the liquid radiator. The mechanical containment for the liquid will have a carbon fiber upstream window and a transparent downstream window made of quartz for radiation resistance.

Eventually a small scale prototype will be built that will contain a liquid radiator and small PMT array, to verify our estimates of Cherenkov photon yield and Cherenkov angle resolution. The beam test of such prototype is planned a year after we test the gas radiator with the HPDs.

7.4.6 Magnetic Shielding of Photodetectors

Both HPDs and PMTs will be exposed to some fringe magnetic field from the return flux of the dipole magnet. Magnetic field calculations predict that the largest magnetic field in the HPD area will be around 100 Gauss; PMTs will see up to 15 Gauss (for the unshielded situation).

We performed calculations of distortions of electron trajectories inside the HPD tube and concluded that we can tolerate up to 5 Gauss field along the axis of the tube, and 0.25 Gauss in transverse direction, after shielding. We plan to verify these calculations by operating the HPD in a magnetic field on our test bench. To shield the magnetic field, the HPDs will be mounted inside a 1mm thick cylindrical mu-metal tubes extending 5 cm beyond the photocathode. We measured a shielding effect of these tubes on a bench and concluded that we can tolerate magnetic fields of up to 30 Gauss in the HPD area. The whole HPD array will be placed in a shielding box (mu-metal and iron sandwich) to reduce the 100 Gauss field down by a factor of 4.

PMTs will also be placed inside mu-metal tubes. These tubes can extend only up to 2 cm beyond the photocathode, otherwise significant shadowing effects occur. We operated shielded 3" PMT from Hamamatsu in the magnetic field and found good resistance to even strong transverse fields. The longitudinal field smoothly deteriorated single photoelectron detection efficiency. For the small fraction of PMTs exposed to the highest magnetic fields we may lose up to 15% of light yield, if all the field is longitudinal.

7.4.7 Radiation Damage Studies

The photon detectors and their readout electronics are situated beyond the aperture of the detector, and therefore are shielded from the interaction point by the dipole magnet elements. Our simulations indicate that the flux of slower particles bent by the magnet onto the PMT array will produce a delivered dose of up to 1 krad/year in the hottest spot. Radiation levels

in the HPD area will be lower by a factor of 20. Measurements done by the other groups show that PMTs with borosilicate glass windows would start deteriorating at an absorbed dose of 100 krad. We plan to conduct our own radiation damage studies for PMT windows and materials that we are considering for gas vessel window (e.g. acrylic).

7.4.8 Mechanical Design

The first iteration of the mechanical design has been completed for all RICH components. Fig. 7.32 illustrates a support scheme for HPDs. A mechanical mock-up of this structure was fabricated. It is being used to refine the design. The RICH beam test will serve as a test of the HPD array with all its elements.

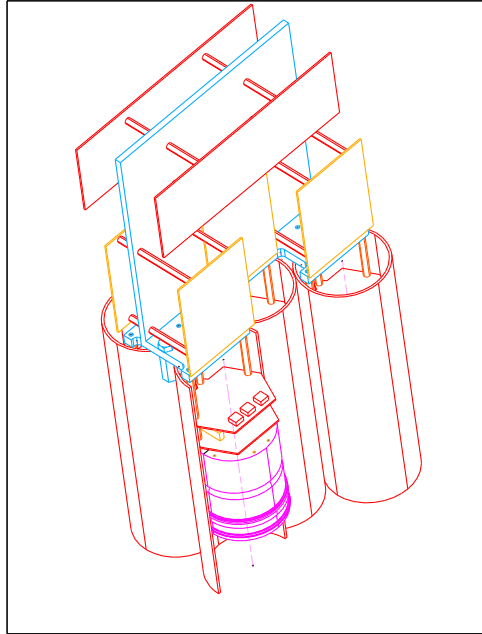


Figure 7.32: Mechanical support of HPDs. They are placed inside of mu-metal magnetic shields. One mechanical unit holds six HPDs (a “hexad”). The “T-spine” supports both the HPDs and the readout electronics.

Mirrors will be supported on three point system mounts that allow for fine adjustments. The prototype support structure was made and it will be used to support a mirror tile in the test beam.

7.4.9 Beam Test

A beam test of the gas radiator, together with a mirror tile and small HPD array is planned for the end of 2002. The test beam has threefold goal: to focus our R&D on HPDs and their readout, to prove that HPDs can be operated in closely packed array under realistic beam

conditions and finally to verify our calculations for the expected light yield and Cherenkov angle resolution.

Progress towards the test beam of HPDs has been already described in the previous sections. Machining of the gas tank for the test beam will start soon at Syracuse.

A separated test beam of liquid radiator with an array of PMTs will be conducted a year later. Since this is not a new technology there is less urgency to perform this test.

7.5 Electromagnetic Calorimeter R&D

7.5.1 Introduction

We discuss here R&D on the three main aspects of the EM calorimeter. The key elements are the PbWO_4 crystals, the mechanical support and the readout electronics. In the final design the crystals will be wrapped with a white reflective paper, be glued to six stage photomultiplier tubes and read out with a modified version of the QIE chip. There will also be a light fiber attached to the crystal so we can input light from an LED.

Thus far we have had two test beam runs at the Institute of High Energy Physics, Protvino, Russia where the crystals have been exposed to both electrons and hadrons. In this case the crystals were wrapped with teflon and coupled to 10 stage Hamamatsu R5800 tubes, read out with a conventional ADC.

We will also discuss the mechanical design for the crystal holder and the electronics.

7.5.2 PbWO_4 Crystals

The absolute light output of PbWO_4 crystals, often referred to as PWO, directly influences the energy and position resolutions. The uniformity of the light output also effects the resolution, especially at high energies. Since the crystals will be exposed to large particle fluxes, they have to be radiation resistant and any light output changes must be monitored. Thus, we need to understand the quality of currently produced crystals, in terms of their total light output, light output uniformity, and radiation hardness.

A test beam facility was constructed at IHEP, Protvino, Russia. It consists of an energy tagged and tunable electron beam and both moderate and ultra-high intensity hadron beams. We studied crystals produced in Shanghai, China, Bogoroditsk, Russia and Beijing, China.

Last year (2001) we have performed beam test studies of the energy and position resolution, effect of longitudinal uniformity on the energy resolution, and radiation damage of the PWO crystals. The crystals were arranged in a 5x5 matrix. The electron beam was directed in the center of the matrix. The beam momentum and position were measured using four sets of drift chambers and an analyzing magnet. To disentangle various contributions to the energy and position resolution, several electron beam energies - 1, 2, 5, 10, 27, and 45 GeV - were used.

The energy resolution is described well using the function $\frac{\sigma_E}{E} = \sqrt{a^2 + \frac{b^2}{E} + \frac{c^2}{E^2}}$ (E in GeV), where a represents a constant term arising from calibration errors, leakage - mostly from the back of the crystals, and non-uniformity in the light collection efficiency along the length of the crystals. The stochastic term, $b = (1.8 \pm 0.1)\%$, arises from photon statistics and leakage of shower, mainly in the transverse directions outside the 5x5 crystal array. The last term $c = (2.4 \pm 0.2)\%$, usually arises from noise of the photon detection electronics, which in our case is negligible. In our studies, the momentum measurement error due to multiple scattering of electrons in the beam line contributes 2.2% to this term and accounts for what we observe. The beam test results for the energy resolution are shown in Figure 7.33.

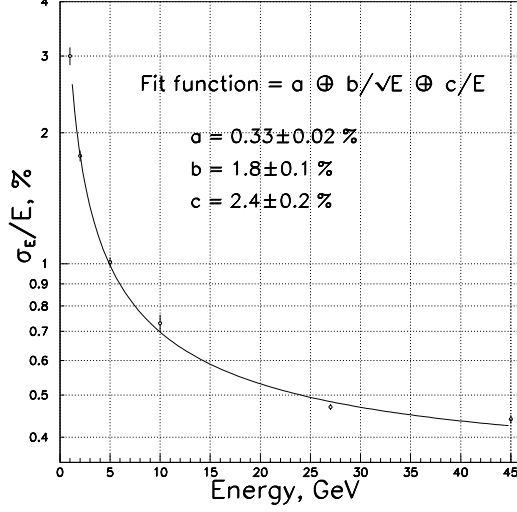


Figure 7.33: Measured energy resolution of the 5x5 crystal matrix.

The constant term is measured to be $a = (0.33 \pm 0.02)\%$. Our Monte Carlo studies show that the shower fluctuations and non-uniformity in light output contribute 0.23% and 0.27%, respectively. The measured longitudinal non-uniformity was used as an input in this Monte Carlo study. The total contribution to the constant term from taking these two terms in quadrature (denoted later as \oplus) is estimated to be 0.35%.

The same Monte Carlo studies show that shower fluctuation results in 0.72% contribution to b . To estimate the other major contribution in the b term we need to know the photo-electron yield. The vendors of the crystals BTCP, Russia and SCI, China measured this number to be about 10 pe/MeV using Cs^{137} and Co^{60} gamma sources and 2" PMT's with bialkalai photo cathode, covering the entire crystal end. Since the PMT's used in the beam test have sensitive areas of (22 ± 1) mm diameter attached to the crystal ends measuring 27 mm square, the covered area is only $(52 \pm 5)\%$. This implies that photo-electron yield in our beam test studies is 5 pe/MeV, and its contribution in the b term is $(1.45 \pm 0.07)\%$. Combining these two contributions as well as an additional contribution from non-uniformity to the b term, we expect b to be $(1.68 \pm 0.07)\%$, which is consistent with the measured value of $(1.8 \pm 0.1)\%$.

One of the major obstacles in our reaching these final resolution results turned out to be Photomultiplier (PMT) gain variations, which result from intensity variations of the electron beam. To monitor short term gain variations light pulses from a LED light source was injected into the crystals through optical fibers. We then read out the LED signal variation as a function of time. The time constant of the variation in PMT output was found to be typically about 10 minutes as shown in Figure 7.34 on the right, where the average pulse heights from LED pulser signals over 90 seconds are shown as a function of time. The left plot in Figure 7.34 shows correlation between the average LED signal and electron beam

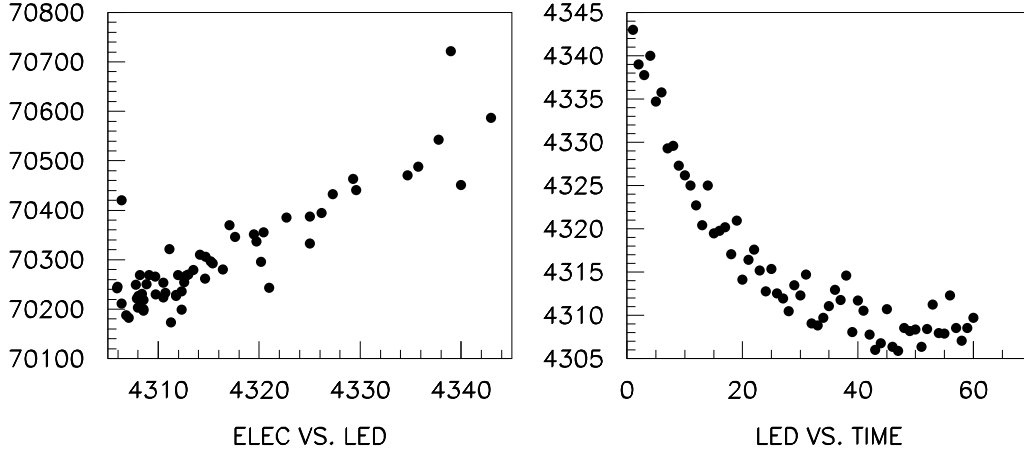


Figure 7.34: Correlation between the LED signal and 5x5 sum for electrons (left). Time dependence of the LED signal (right). Each point is an average over 90 seconds. Note the highly suppressed zero on the scales. (Data are at 27 GeV.)

pulse height. The correlation is good and the slope is close to 1. When the LED signal is used to correct the beam pulse heights, we were able to obtain the expected resolution.

The position resolution of the 5x5 matrix was also obtained using the test beam data. The resolution, averaged over electrons spread across the entire central crystal was calculated for each beam energy. From a two-term fit to the position resolution data we obtain $\sigma_x = (0.28 \pm 0.008) \oplus (3.32 \pm 0.02)/\sqrt{E}$ mm (E in units of GeV). It agrees well with the resolution expected from Monte Carlo simulation, which is $\sigma_x = 0.2 \oplus 3.5/\sqrt{E}$ mm.

The crystal longitudinal light output uniformity was measured using both radioactive sources at Minnesota, Syracuse and Protvino, and using a transverse muon beam at Protvino. The data were binned into 1 cm bins along the crystal length. For each bin the pulse-height distribution was fitted to a modified Landau distribution to obtain the peak position. The peak positions were fitted as a function of the coordinate along the crystal. For a group of the 25 crystals we obtain an average slope of 0.55%/cm with an rms variation of 0.15% in the region of the shower maximum 3 – 10 radiation lengths from the front of the crystal. We do not observe any significant differences between the 5 Bogoroditsk and 20 Shanghai crystals. These measurements were consistent with the source measurements, but much more precise due to the high light yields.

To study the radiation hardness of the crystals, we irradiated them with intensive pion and electron beams using dose rates of up to 60 and 30 rad/hour, respectively. A few crystals were exposed to much higher radiation of 100 krad/hour in a dedicated facility for up to 2 Mrad of accumulated radiation (only crystals were exposed in this case, not the PMTs). Most crystals which received a dose rate of 30 to 60 rad/hour lost about 20% of their light

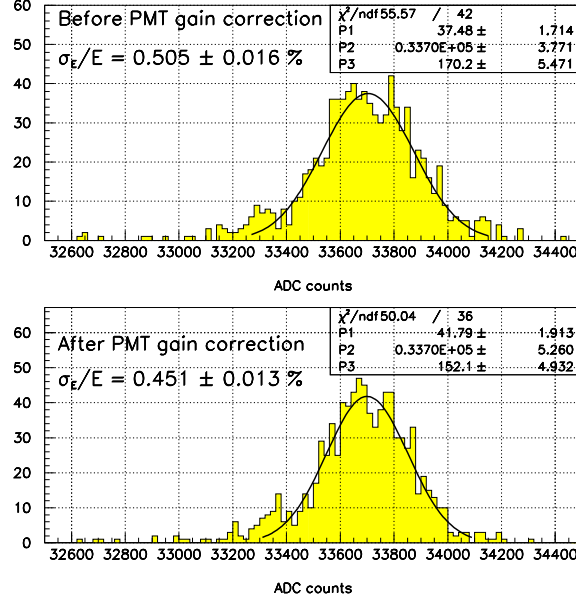


Figure 7.35: Energy resolution before and after the PMT gain correction using the LED pulser information. Data are collected using a 27-GeV electron beam.

outputs. The remaining crystals received up to 10 rad/hour and lost between a few per cent and 20%. With the electron beam, the light loss was about 15 to 20%. The small difference between the damage due to electron and pion radiation appears to be mostly accounted for by their difference in the radiation profile along the length of the crystal. With the electron beam, the radiation level reaches the highest level 5-7 cm into the crystal and falls sharply by two orders of magnitude at the PMT end, whereas with the pion beam, the radiation level stays more or less constant (within a factor of 2) after the maximum. A small effect due to the difference between the physical processes by which electrons and pions interact with matter cannot be ruled out.

At the extremely high radiation rates, which is 100 times larger than the highest radiation dose rate that any crystals in the BTeV environment will suffer, the two crystals, one from Bogoroditsk and the other from Shanghai, lost 2/3 of their light. These crystals received close to the maximum integrated dose for the worst case in BTeV in a few days, and there was no time for the natural recovery process in the crystals to work. (Typically, the crystals recover from radiation faster than the radiation rate in BTeV.) It is really reassuring that these crystals survived even this extreme environment.

Various tests of the PMT's proved that the light loss measured above was not due to radiation damage of the PMT's we used, which had regular borosilicate glass windows. The total dose they received was less than a few krad. These tests include subjecting the same PMT's to over 10 krad of radiation using a Cs^{133} source at Minsk, and a similar test done

at Minnesota using only window glasses obtained from the PMT manufacturer. However, during the fall 2001 beam test run, we found that the gain of PMT's easily changes by a few per cent and sometimes up to 5% depending on the history of signal intensities. These studies were initiated when we observed a small pulse-height change in the LED signal when the accelerator went down for a few hours. On a test bench, we analysed pulse heights when different DC light was superposed to pulsed light. We observed that the pulse height changes a few per cent even when the DC light produced less than $1\text{ }\mu\text{A}$ of anode current, and when the average anode current is increased to $10\text{ }\mu\text{A}$, the change did not increase very much but it became more permanent. The change had a time constant of order 10 minutes.

From these observation, we had to conclude that non-trivial part of the light loss, which we attributed to radiation damage of crystals above, might be due to PMT gain changes caused by signal rate changes since the radiation dose level was controled by the signal particle intensity. Therefore, the numbers quoted above must be understood as upper limits.

We are working on the data analysis, where we are trying to separate the two effects: the PMT gain change and crystal radiation damage, using their differences in the time constants of the changes. Since the changes appear to depend on the history and not easily characterizable, if we want to know the radiation damage effect at a per cent level, we may have to change our strategy. We plan to use red LED light pulser (in addition to the current blue one) to monitor the PMT gain change alone in the future beam tests. Transmission of red light in the crystals is not effected by radiation damage.

We plan to continue beam tests with more crystals purchased from Shangai, Bogoroditsk, Beijing and possibly another facility at Apertiti, Russia. We need to evaluate the characteristics of PWO crystals, which will allow us to decide on crystal specifications which balance the physics performance and the cost of the calorimeter.

7.5.3 Mechanical Support Structure

The crystal support structure was assumed to be similar to the CMS endcap calorimeter for the 2000 proposal cost estimate, but we considered the \$2.2M cost as rather expensive. In the last year, we have explored a more conventional design based on a grid structure using interlocking sets of vertical and horizontal aluminum strips instead of carbon-fiber cell-structure-based "CMS" design in an effort to bring down this cost, while maintaining the excellent calorimeter resolution.

The new design is depicted in Fig. 7.36, where, columns of crystals will be supported from 0.3 mm thick aluminum strips (or similar thickness carbon fiber) shown as closely spaced vertical lines in the Figure, that are strung vertically from top to bottom on a large frame structure. There will be a vertical strip between each column of crystals. Crystals will rest on 0.3 mm thick horizontal strips (closely spaced horizontal lines drawn only in the upper right quadrant) which run though slots machined in the vertical strips. There will be two horizontal strips (front and back) for each row of crystals. The horizontal strips will be anchored to the sides of the frame structure. The load of each crystal (1.4 kg) is transferred from the two horizontal strips to the adjacent vertical strips. The gap between the edge

of the crystal and vertical strip is small enough that the horizontal strips do not deflect significantly. The transfer of crystal load to the vertical strips is highly distributed as each pair of horizontal strips transfers the load of one crystal. The vertical strips transfer the load of a column of crystals to a mounting block over a large shear area. The mounting blocks attach to the frame structure with bolts.

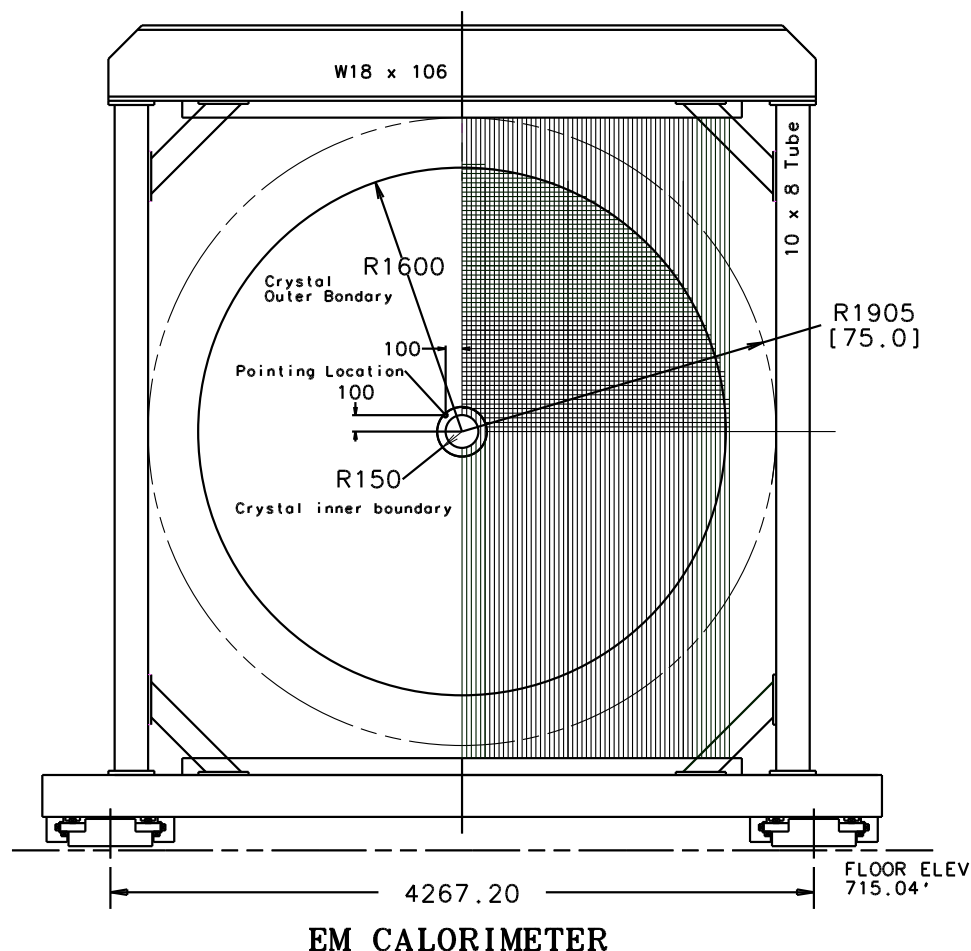


Figure 7.36: Front view of the EM calorimeter support structure. Dimensions are in mm except as indicated.

The frame is a simple square shaped structure. The top is a stiff beam that transfers the load of 100+ vertical strips to the two side columns. The base of the frame structure is attached to bearings and rails that allow the entire structure to move smoothly along the beam direction. The load paths of the structure are statically determinate and allow for design with relatively simple analysis. The stress in the most highly loaded vertical strips will be about 40% of the yield strength of sheet aluminum at the 2g loading condition. The vertical strips are oriented at the proper angle by pins located in the frame structure. Machining the slots in the vertical strips at the proper angle orients the horizontal strips. The

spacing of the vertical strips will be maintained along the entire length by small interlocking tabs and slots in the vertical and horizontal strips.

The cost effectiveness of the structural design depends on the ability to economically machine the thousands of slots in the vertical and horizontal strips. Laser cutting, which is a common commercial process, will be used to machine the slots. To evaluate the design concept and to obtain a preliminary cost estimates of cost-driving elements a prototype with 4 vertical and 12 horizontal columns is being constructed. The vertical strips for the prototype are 1/2 the length of the final design. Pictures can be viewed at http://home.fnal.gov/~howell/ecal_prototype_photos/ECAL_prototype_details.html.

7.5.4 Electronics

The PMT bases are being developed. A schematic diagram has been created and layout work is beginning. We plan to use a version of QIE for the front-end electronics, and have talked to the designer, as well as a few electronics engineers who have designed supporting electronics for previous versions of QIE. This chip was used in the KTeV calorimeter and we have confirmed that QIE is a good match to our needs.

When PPD engineers become available we will start these activities:

- A BTeV version of the QIE chip will be developed along with test boards that will serve in the process of evaluating the performance of the chip.
- A multi-channel ADC circuit board that will include the QIE chips will be developed.
- A PMT Cable Transition Card will be used to fan-in the signals from the PMT bases to the ADC cards.
- A Controller/DCB card will be used to interface between the ADC cards in a crate and the DAQ system.
- A commercial computer processor card will be selected to interface between the control system and the crate. A custom card will be developed to convert between the form factor of the commercial processor card and the form factor of the crate.
- A custom J3 backplane will be developed to facilitate I/O operations between the Controller/DCB card and the ADC cards.

7.6 BTeV Muon System R&D

7.6.1 Introduction

The BTeV muon group consists of three institutions, Vanderbilt University (lead institution) (VU), University of Illinois at Urbana-Champaign (UIUC) and University of Puerto Rico at Mayaguez (UPR). Research and development work on the muon system started in earnest when the group decided on the detector design and technology in 1998/99. Since then, the group has made significant progress on determining how the detector will be constructed and supported as well as a better understanding of the electronics layout.

7.6.2 Summary of the BTeV muon system

7.6.2.1 Baseline geometry

As shown in Fig. 7.37, two toroids, each 1 m long with 1.5 T fields provide the bending power for determining the muon momentum. The muon detectors will be set up in three stations, one between the toroids and two behind the toroids. The momentum can be measured using the two, well shielded, downstream stations and the nominal beam constraint. The station between the two toroids provides a powerful confirming hit to eliminate fake tracks.

7.6.2.2 Baseline detector

The basic building block in the construction of a detector station is a “plank” of thin walled (0.01”) 3/8” diameter stainless steel proportional tubes as shown in Fig. 7.37. Thirty-two tubes, arranged in a double layer with an offset of half a tube are soldered at each end to a brass manifold and supported in the middle by a brass rib piece. This provides a sturdy, self-supporting building block which also acts as a Faraday cage to reduce external RF noise. Proportional tubes were selected because they are robust and have the necessary rate capability.

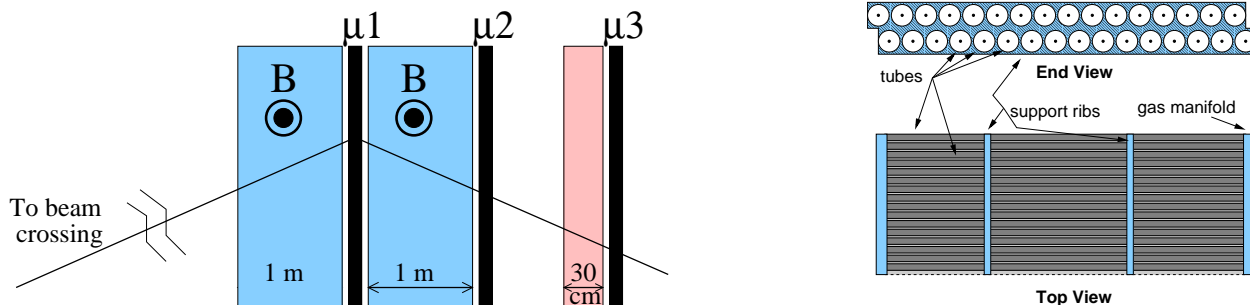


Figure 7.37: (left) Muon system geometry. (right) Views of one plank of proportional tubes.

The 0.5 cm wire spacing of this design has an effective spatial resolution of $5 \text{ mm}/\sqrt{12} = 1.4 \text{ mm}$ with no dead regions between tubes. This meets our requirements for momentum resolution, drift time and occupancy.

To minimize occupancy at small radii and improve pattern recognition, each detector station consists of eight overlapping pie shaped “octants,” shown in Fig. 7.38. There are four views (r , u , v , and r) in each octant as shown in Fig. 7.38. The r (radial) view is repeated to provide redundancy for the most important view and help reject fake tracks in the trigger. The u and v views are rotated $\pm 22.5^\circ$ to measure ϕ and resolve hit ambiguities, reducing the misidentification rate. The views stack on top of each other and are built from 12 planks/view. Pairs of octants, combined into quads, will be the structure moved into the spectrometer.

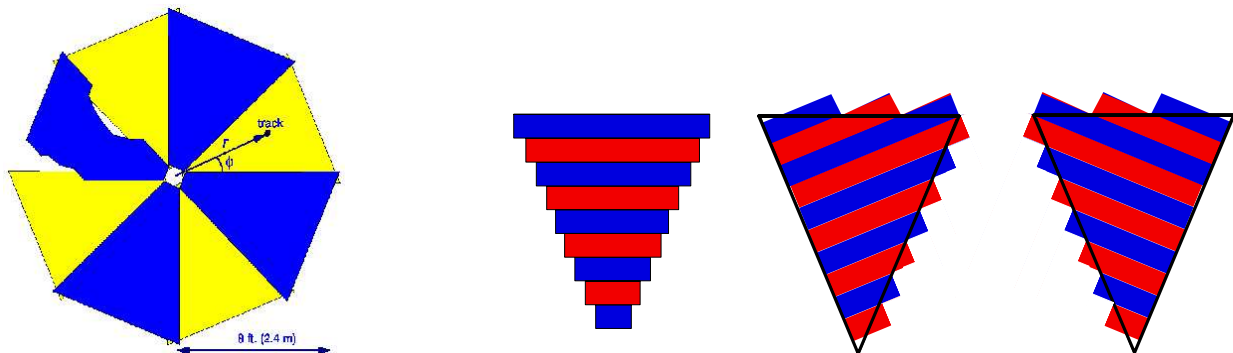


Figure 7.38: (left) Beams-eye view of one muon station (eight overlapping octants arranged in two layers). (right) Arrangement of planks to form each of the four views in an octant (r view is repeated). There will be 12 planks per octant (more than shown).

7.6.2.3 Front-end electronics

The front-end electronics will be similar to that used for the CDF central outer tracker (COT): one circuit board to deliver high voltage and another circuit board with electronics to amplify and digitize the tube signal. Both boards will be located directly at the end of a group of 16–32 proportional tubes.

We plan to use the ASDQ integrated circuit developed at the University of Pennsylvania to amplify and digitize the signals coming from the proportional tubes. This chip is used in the Run-II CDF COT for a similar purpose. The ASDQ, amplifies the first ~ 8 – 10 ns of the signal and outputs an LVDS signal. This chip, when mounted on a circuit board, has a low effective threshold of about ~ 2 fC (confirmed by tests at VU). The chip also features a double pulse resolution of ~ 20 ns. The ASDQ digital signals will be sparsified, serialized, and read out using a standard Fermilab readout protocol. Fiber optic cables will transfer the data from a combination of 12 planks to a buffer memory.

7.6.3 Past research and development work

7.6.3.1 Detector design

The initial job of the muon group was to determine the detector design and technology to be used. The goals of the muon system were twofold:

1. Operate a dimuon trigger, independent of all other detectors, to provide calibration for the more ambitious vertex trigger and enhance the physics reach of dimuon events
2. Provide muon identification

Initial calculations and Monte Carlo studies were performed by the University of Illinois to obtain a rough idea of the detector requirements which would satisfy the muon system goals. These studies were used to set the initial toroid requirements, detector resolution, and station separations. The choice of a proportional tube system was influenced by many considerations including cost, neutron background, and the University of Illinois experience in building muon proportional tube systems.

7.6.3.2 Plank design/construction

The first round of plank prototypes (10 planks of 32 tubes each) were constructed in the first half of 1999. These planks were constructed in the following way:¹

1. Tubes were cut to length in the machine shop from purchased stock.
2. Each tube was cleaned in an Alconox solution, rinsed, and dried with compressed air.
3. As shown in Fig. 7.39, a gold-plated tungsten wire was strung through a Delrin endpiece on one end of the tube, through the tube, and through another Delrin endpiece. The Delrin endpiece consists of a circular piece of Delrin with a lip at one end to hold it at the edge of a tube. A hole drilled through the center of the Delrin contains a small brass tube (crimp pin) extending out. The brass crimp pin also contained a double funnel inside to center the wire. Each Delrin endpiece also had three small holes for gas flow.
4. After stringing, one end of the tube was “crimped.” The Delrin piece was inserted into the tube. Then a resistor lead was inserted in the brass crimp pin (along with the wire) and a commercial crimp tool was used to crimp everything together.
5. After crimping one end, the other end was attached to a calibrated weight to achieve the proper tension and the other end was crimped.
6. Continuity and high voltage tests on each tube ensured the crimp held and the wire did not break.

¹pictures at http://www.hep.vanderbilt.edu/~sheldon/tubes/tube_fab.html

7. A plank was constructed from 32 strung tubes. The endcaps were machined from Noryl (plastic) and contained one hole for gas and 32 small holes for the end of the crimp pins (which connect to the electronics). The endcaps were glued to the end of planks.

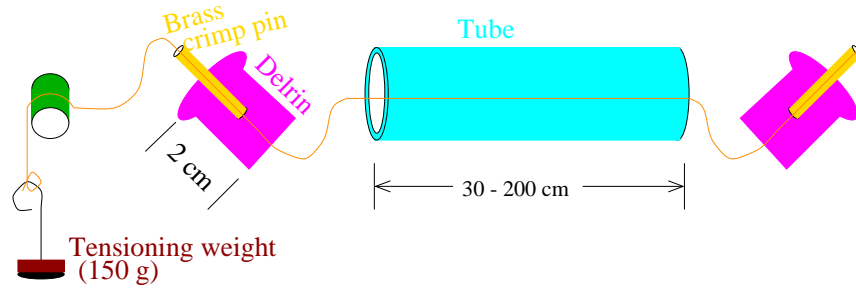


Figure 7.39: Diagram of the stringing process.

The construction of 10 planks in the spring of 1999 at VU provided us with valuable information. We found that $30\text{ }\mu\text{m}$ and $50\text{ }\mu\text{m}$ wire both work well while the $20\text{ }\mu\text{m}$ wire was harder to string and was not needed for the muon system. We found that one crimp often did not hold the wire in place while two crimps were almost always sufficient. The crimp was also not airtight requiring a glue or solder on the end of the crimp pins to ensure a good seal.

The 10 planks (six 1-footers, two 3-footers, and two 6-footers) were transported to Fermilab in June, 1999 for a test beam run. A stand on which the planks could be mounted was also designed and built at Vanderbilt which allowed the planks to be rotated and offset.

The front-end electronics were sample ASD8B cards from the University of Pennsylvania. Testing these boards revealed a high susceptibility to ambient RF noise. To reduce this noise, boxes to enclose the electronics were constructed out of circuit board and wrapped in copper tape. Interface cards to provide high voltage to the tubes and low voltage to the electronics were also designed and assembled. An interface card to convert the LVDS signal to ECL was also designed and built. High voltage, low voltage, and signal cables were also assembled. This work was done by VU and UPR.

The muon data acquisition system was written by VU using a CAMAC interface. TDC data from the planks and latches from the trigger scintillators were recorded. UIUC wrote much of the reconstruction software and an online event display. Participants from all three institutions took shifts for approximately 16 out of 24 hours for most of August and September. Pictures of the test beam run can be found at http://www.hep.vanderbilt.edu/sheldon/tubes/test_beam.html.

Plateau curves on the planks showed that the planks are $>95\%$ efficient at 1.8 kV for a $30\text{ }\mu\text{m}$ wire with Ar/CO_2 gas which met our expectations. However, cross talk between channels was very high which resulted in many tubes in a plank firing with only one incident particle. This severely hindered attempts to measure individual tube efficiency or position resolution.

The susceptibility to external noise and extensive cross talk led to several design changes. One of the successors to the ASD8B chip (used in the test beam) was the ASDQ (proposed for the real muon system). Tests at VU proved that this chip was more resistant to external RF noise. We also redesigned the high voltage distribution card. To further protect the electronics from RF noise, plank design was changed to use a brass endpiece (instead of plastic) which was soldered to the stainless steel tubes. These two modifications solved the cross talk and external noise problems. The changes to the design also required changes in plank construction. Since the Delrin piece might melt or slip during the soldering process, the tubes needed to be strung after soldering the endpiece. A new new homemade crimp tool was created at VU which works in the restricted space available. The other change made was an aluminum box which contains the electronics and bolts to the endpiece, providing the last part of the Faraday cage. The Delrin piece and brass crimp pin design were also redesigned and supplied by UIUC. Pictures of the new design can be found at <http://www.hep.vanderbilt.edu/~wjohns/res/btev/proto2/>.

7.6.3.3 Front-end design

At VU and Fermilab, work continues on developing the board which will be partnered to the ASDQ. The board will take 16 or 32 channels from a plank and send the data out. Initially, an analog-only board was proposed by VU; similar to the CDF (Run II) central outer tracker board. The proposed board would send LVDS signal through ribbon cable to an interceptor board which would digitize and timestamp the pulse and feed the data through a serializer and into the DAQ system. Fermilab proposed an alternative solution which combines these two functions into one board. Development work is continuing on this combined board. A prototype has been built at Fermilab and appears to satisfy the muon system requirements. Further testing of this board will take place in the summer, 2002 beam test. Initial testing results can be found at <http://www.hep.vanderbilt.edu/~wjohns/res/btev/proto2a/>.

7.6.3.4 Detector construction/support

UIUC is developing detailed mechanical designs for mounting and assembly of the BTeV muon planks into quadrants. This involves considerable drafting work as well as the construction of a detailed 1/5-scale model of a complete quadrant. The scale model will help us formulate a plan for installation of the muon system as well as a plan for replacement and repair of broken planks during the run.

UIUC has also designed an adjustable jig with controllable heating elements in order to maintain precision tube lengths and properly orient the end caps with respect to the tubes and mounting plates during the assembly process. Such a jig is required to efficiently assemble the large number of tubes of many different lengths with adequate precision to allow us to mount the array in the confined space available.

7.6.3.5 Simulation work

All three muon group institutions have been involved in creating an accurate muon simulation in the BTeVGeant framework, including the complicated magnetic field resulting from having a dipole magnet inside a toroidal magnet (see chapter 4 for a discussion).

7.6.3.6 Construction database

At each step in the construction process the actions taken for each element of the detector are recorded in a database through the use of a barcode tracking system. We anticipate that these data will later be useful in inventory control and studying correlations of detector performance parameters.

7.6.3.7 Wire tension measurement

To assure that each tube is strung with a wire of the correct tension, the tension of each wire is measured by placing the plank in a magnetic field, driving the wire with a sinusoidal current, and measuring the induced EMF to find the resonant frequency. From this value the tension is computed.

Vanderbilt has developed a test stand which automatically measures the tensions in a plank full of tubes and stores the resulting information in the construction tracking database.

7.6.4 Future research and development work

During the past year we have constructed three planks at VU using our latest design, worked on by all three institutions. These planks are performing well in a cosmic ray test stand as shown in <http://www.hep.vanderbilt.edu/~wjohns/res/btev/proto3/>. We plan to construct 3–5 more planks to test during the summer of 2002. This test run will allow us to make many of the studies which we were unable to perform in 1999 due to noise problems. These studies include measuring speeds and responses of various gasses, measuring high rate effects, and measuring individual tube efficiency and resolution. The results of these tests should verify our design will be effective in the BTeV spectrometer.

Refinements of the plank construction will also be investigated in hopes of finding a method which reduces the construction time while maintaining the needed features. In particular, it has been proposed that tubes may be strung before joining the tubes to the endpiece. The problem of heating the Delrin piece beyond its melting point may be resolved by spot welding the tube and endpiece at a single point and filling the rest with glue rather than soldering the entire structure.

Work on the front-end electronics continues at VU and Fermilab. We plan to test the new combined analog/digital board during the summer, 2002 test beam.

UIUC will continue work designing tools to assemble and support the muon system.

All three institutions will also continue working on simulations and start writing reconstruction code.

7.7 Trigger R&D

One of the key features of BTeV is the vertex trigger, which is the primary trigger for the experiment. Most of the trigger R&D work prior to the submission of the BTeV proposal has been devoted to the vertex trigger, especially the implementation at Level 1. Since the submission of the proposal we have made substantial progress in the development of the Level 1 vertex trigger. We have also included the Level 1 muon trigger in our hardware design, made substantial improvements to the muon trigger algorithm, and we have performed a large number of studies as we continue the development of Level 2 and Level 3 trigger algorithms. Progress on all of these R&D developments will be presented in this section.

7.7.1 Level 1 Vertex Trigger

The Level 1 vertex trigger continues to be an important aspect of our trigger R&D effort. Since the submission of the BTeV proposal, this R&D work has focused on the following areas:

- Implementation of the pattern recognition (segment finding) algorithm in FPGA's (Field Programmable Gate Arrays)
- DSP (digital signal processor) timing studies for the segment matching, tracking, and vertexing portion of the Level 1 vertex trigger
- Timing studies on processors other than DSP's
- Data flow analysis and simulations
- Design and fabrication of a 4-DSP prototype board

We have new results on DSP timings that were obtained by optimizing the Level 1 vertex code for a DSP, and then running the code on a DSP. This approach gives us a more reliable assessment of the number of DSPs needed for the vertex trigger, as opposed to the approach we used for the BTeV proposal, for which we estimated CPU cycle counts based on a partial assembly language implementation of the code. With the optimizations and timing results achieved to date, the DSP code is close to the timing estimates presented in the proposal but a factor of 4 slower. This is reduced to a factor of 3 for a new DSP that will be available beginning in the second quarter of 2002. Even without further optimizations, we are confident that newer generations of DSP's will get us to the estimates stated in the proposal. However, in an effort to explore other processor alternatives for the Level 1 trigger we have started to investigate general purpose processors from Intel and Motorola. Preliminary results suggest that these processors exceed the proposal estimates by at least a factor of 2 or 3 in performance. Additional R&D will be needed to investigate the feasibility of implementing the Level 1 trigger hardware with these processors.

7.7.1.1 DSP Timing Studies

This section will focus on timing studies done on the segment matching, tracking, and vertex finding portion of the Level 1 vertex trigger. Since this is the portion of the Level 1 vertex trigger algorithm that will run on the DSPs, it will be referred to as the DSP algorithm. We stress the importance of timing results since execution speed directly determines the total number of DSPs required in the Level 1 farm and therefore its cost and complexity. We will also describe work done to address two concerns raised by the PAC's Technical Review Committee in 2000, specifically (1) the fact that the custom assembly language code described in the proposal had not been run on a DSP or DSP simulator to demonstrate that it works, and (2) the potential problems involved in maintaining custom assembly code throughout the lifetime of the experiment.

The starting point of our studies was a variant of original C language version of the code on which the optimized assembly language version used in the proposal was based. (The code whose performance we describe here was optimized for triggering efficiency. The code used in the proposal made small sacrifices in triggering efficiency to achieve significant enhancements in computing speed. This will be discussed more below.) All DSP timing studies were done on a Texas Instruments (TI) C6711 DSP Starter Kit (DSK) board with a 150MHz TI TMS320C6711 DSP.

The initial step taken was to get the original, unoptimized C code running on the DSK board to address issue (1) above. Once the code was running on the DSK, optimizations were introduced in two major phases to reduce the execution times. In the first phase, costly calls to external library functions were replaced with C intrinsics that map directly to DSP instructions, and all double-precision floating point operations were replaced with single-precision operations. In the second phase, the segment matching portion of the algorithm which accounted for over half of the total execution time was completely rewritten in order to avoid the need to try all possible combinations of inner and outer segments in searching for complete tracks. Data structures of the code were also reduced in size to allow assignment of data memory sections in the DSP's internal Level 2 Cache/SRAM. This reduces CPU pipeline stalls due to cache misses requiring fetches from external SDRAM. In order to address issue (2) above, we also made it a point to do all optimizations in C, resorting to assembly language programming only when necessary.

Execution times after each phase of optimizations were measured using the built-in profiler of TI's Code Composer Studio to count CPU cycles accumulated in the DSP's on-chip performance monitors. Due to the reasons cited above, execution times are improved by having code and data memory sections reside in internal RAM. Since the DSP code² exceeded the 64KByte size of the internal memory, groups of functions were profiled in separate sessions wherein only those functions being profiled were assigned in internal RAM. This was done with the view that upcoming processors in TI's roadmap will have enough internal memory to hold the complete DSP code.³ The measured execution times for both optimiza-

²roughly 2,500 lines of C source code representing on the order of ~ 100 KBytes in .data and .text memory sections.

³A new member of C671x family-TMS320C6713 will be sampling in 2Q02 with 256KBytes of on-chip L2

Section of DSP algorithm	Proposal estimates in CPU cycles	DSP timing in CPU cycles per beam crossing			
		Before opt. 10 BCO's	After Phase 1 opt. 10 BCO's	After Phase 2 opt.	
				10 BCO's	100 BCO's
Segment matching	24,200	1,296,778	503,012	164,836	168,113
Track processing	14,400	397,518	38,632	38,632	34,528
Vertexing	14,673	264,429	34,720	34,720	32,938
Total	53,273	1,958,725	576,364	238,188	235,579

Table 7.3: Level 1 Vertex Trigger Timing Results

tion phases are presented in Table 7.3 where they are compared with estimates in the BTeV proposal. These results were obtained with 10 simulated (Pythia/Geant3) minimum-bias beam crossings with an average 2 interactions per crossing. For verification, the timing results for the second phase of optimizations was repeated with 10 times more statistics (Table 7.3, column 2). As these results show, execution times were reduced by nearly an order of magnitude ($1,958,725 \text{ cycles} \rightarrow 235,579 \text{ cycles}$) without resorting to custom assembly programming.

7.7.1.2 Optimizations with Custom Assembly Code

To see what performance gains could be achieved with custom assembly programming, code coverage of the segment matching routine was tested to identify sections on which optimizations would have the greatest impact. Using the intermediate assembler source code (totalling 1,329 lines) produced by the compiler for this routine as a starting point, custom assembly programming was done on the identified sections to make more efficient use of the DSP's 8 parallel execution units. 222 new lines of custom assembly code were written replacing roughly 20% of the compiler generated assembly code. Applying the procedure described above for measuring execution times on this optimized assembly version of the segment matching routine indicated a reduction of $\sim 2\times$ for this particular routine ($117,242 \text{ cycles} \rightarrow 63,618 \text{ cycles}$) and an overall reduction of over 20% for the complete DSP algorithm ($235,579 \text{ cycles} \rightarrow 181,955 \text{ cycles}$).

7.7.1.3 Efficiency-Speed Tradeoffs

As mentioned above, the code used at the time of the proposal made some small sacrifices of ultimate triggering efficiency to achieve speed ups which explain in part the differences reported above. The version of the algorithm used here is aimed at maximizing the triggering efficiency. We are now going back and reviewing those tradeoffs. However, given the likely availability of faster DSPs, we do not see this as an urgent priority. Moreover, as will be seen below, alternatives to DSPs are now available which may already have provided us with a

SRAM (enough to hold the complete DSP algorithm) and higher core speeds of 225MHz (33% reduction in execution times).

new approach which resolves nearly all remaining issues with respect to the Level 1 trigger speed.

7.7.1.4 Timing Studies on Other Processors

Investigations have also been done to measure execution times of the optimized C version of the DSP code on the following two general purpose processors: (a) 1.13 GHz Intel Pentium III-M, and (b) 500 MHz Motorola MPC7400 PowerPC G4. Additional optimizations done in C to utilize both processor's built-in vector execution units were introduced in the segment matching routine. Preliminary results are very encouraging indicating execution times roughly an order of magnitude better (Pentium III: $133,053 \text{ cycles} \times 0.88 \text{ ns} = 118 \mu\text{s}$, PPC G4: $98,365 \text{ cycles} \times 2.0 \text{ ns} = 197 \mu\text{s}$) than those for the optimized C code running on a 150 MHz TI TMS320C6711 DSP ($235,579 \text{ cycles} \times 6.67 \text{ ns} = 1,571 \mu\text{s}$). Price/performance ratios of a PowerPC G4 based system seem quite promising compared to that based on a DSP. More work needs to be done, however, to investigate whether the I/O capabilities of these processors meet the requirements of the trigger.

7.7.2 Level 1 Muon Trigger

7.7.2.1 Code Development Infrastructure

We have developed an event driven analysis infrastructure that reads a user-specified list of muon Monte-Carlo input files, translates the event by event raw hit information into a more easily useable set of data structures, and provides the hooks necessary for the user to easily link their own analysis code. Developed in Visual C++, this system runs on any PC and is the basis of all algorithm development done so far.

7.7.2.2 Muon Trigger Algorithm Development and Testing

Work done prior to the submission of the BTeV Proposal in May, 2000 had demonstrated that when all available information is used (i.e. a fitting algorithm is applied to all possible hit combinations and the resulting χ^2 examined), good efficiency for muon tracks and adequate rejection of background could be obtained. The goal since this very important proof of principle has been to see if a simpler (hence potentially much faster) algorithm could achieve similar results.

By using a large set of Monte-Carlo generated “good muons” and studying both the correlations between hits in different views within a station, and the correlations between hits in similar views in different stations, we discovered that some simple relationships exist between hit muon tubes that belong to a single track. Not surprisingly, there is a very tight correlation between the hits in the various views of a given station, and this correlation can be used to distinguish sets of tubes that belong to the same track (which define a space-point in that station) from the otherwise enormous set of random combinations.

Slightly less intuitive is the fact that there is a very simple relationship between the radial coordinates of the three space-points (one per station) belonging to the same track. If we simply let these radial coordinates define points in the 3D vector space r_1, r_2, r_3 , we find that good muon tracks result in points that populate a simple plane in this space, and the random background combinations results in points that tend to not lie in this plane. This allows us to use simple geometric transformations on points in this space to, in effect, construct look-up tables to identify good tracks. Using such an approach, we have designed an algorithm whose performance seems comparable to the full-blown χ^2 fitting method, yet is simple enough that it should be able to run quickly on a DSP.

7.7.2.3 Muon Trigger Algorithm Timing Studies

We have implemented a tube-based lookup oriented trigger algorithm, as outlined above, in C code that runs on a TI TMS320C6711 DSP starter kit. We use this hardware, together with the TI's Code Composer Studio, to study the speed of the algorithm as various modifications are made. The code has not yet been optimized, and currently its speed is about an order of magnitude away from the performance assumed in the proposal. Part of the needed speed improvement will come from simply buying faster devices; part will come from further code optimization, and part will come from moving the space-point finding portion of the algorithm from DSP based code to upstream FPGAs. This is a work in progress, and we hope to have more reliable results in a few months.

7.7.3 Level 1 Hardware

7.7.3.1 Simulations and Data Flow analysis

The lack of queuing models and data-flow analyses of the BTeV trigger were ranked with high technical risk by the PAC's Technical Review Committee in 2000. A data flow analysis is now underway using two different approaches: modeling and simulation. This analysis was carried out through 2001, and will continue in 2002. We follow an iterative procedure: a trigger architecture is proposed, mathematical models are used to analyze the data flow and buffering needs of the system, then simulations are used to validate the model assumptions, and the results are used to refine the proposed architecture. The data flow analysis was extended beyond the Level 1 trigger boundary into the pixel detector readout. The purpose of this is to understand how the readout function scrambles the data with respect to time and to explore the issues involved with time-demultiplexing the pixel data into parallel streams called "highways." Multiple highways reduce the bandwidth per data stream, and the number of highways will be chosen to match the most appropriate technology.

With respect to the Level 1 pixel trigger, the pixel preprocessors and segment finders have been fully modeled and simulated. The data flow analysis shows a very good degree of consistency between the modeling and the simulation results. Queue sizes and communication channel bandwidths are within reasonable margins. Once we begin the design stage and a full VHDL code has been generated, we can refine some of the parameters shown in the

modeling and simulation to optimize the design. In particular, the segment-finder FPGA algorithm is composed of similar module functions that work sequentially. The data flow analysis has helped to optimize the data pipelining and to detect underutilized modules. Some modules can be reused for more than one function. This is particularly important to minimize the silicon logic needed in the implementation.

In 2002 we plan to complete the data flow simulations. This requires modeling and simulating a farmlet of processors, which is the smallest grouping of processors on a single board. In our current design, a farmlet consists of four DSPs. The data flow simulations will be used to study and optimize the data flow from input buffers to the processors. Subsequent simulations will be used to study the dynamics of control and monitoring for the farmlet.

7.7.3.2 Segment Finding Algorithm Implementation and Simulation Status

The segment finding algorithm is being implemented in Very High Speed Integrated Circuit Hardware Description Language (VHDL). We are using this high level language in an attempt to keep the code somewhat portable between the various software tools that are currently available for field programmable gate arrays (FPGAs) and complex programmable logic devices (CPLDs). However, to make the most efficient use of device specific resources such as internal RAM or other custom features, the manufacturer's device specific library components are also incorporated into the design. The segment finding algorithm is currently being executed and simulated using the Quartus II design software from Altera. The simulations have validated the segment finding algorithm and produced two different implementations of the algorithm and may help us to optimize the trade-offs between cost and speed. There have been preliminary compilation, place-and-route, and simulation tests run so that silicon usage and timing estimates can be accumulated from these experiments for future comparisons. These estimates were used as inputs to WBS costs. The next goal is to move the existing VHDL code to other design software platforms. Our intent is to target devices from other programmable logic manufacturers. This will allow us to make direct price versus performance comparisons between device offerings from competing manufacturers.

7.7.3.3 Hardware investigations

There have been ongoing refinements of the Level 1 vertex trigger architecture. A trigger DSP prototype has been designed based on the current architecture. It is implemented as a motherboard containing I/O controllers and four daughter-card positions for processing elements. The daughter-card positions have a generic set of signal assignments to allow various processors to be accommodated. The current target processors are DSPs, and two daughter card designs for specific DSP chips are included in the prototype project.

The motherboard implements the following functions:

1. Input data buffering and management of data flow to the processing elements.
2. Output data buffering and management of data flow from the processing elements to the external data network.

3. Processing of control messages from the external control network.
4. Control and management of result messages generated by each processor.

The daughter board implements the following functions:

1. Contains flash and RAM memory, power supply and support chips for each processor.
2. Provides external power input connectors and serial port connections so that it can be operated stand-alone without the motherboard.

The project goals include:

1. Provide a platform to characterize processor-to-network interactions with several types of processors. This includes testing several data flow and buffer management techniques for the current architecture, which is based on a single data stream feeding multiple processors.
2. Characterize the impact of message traffic on processing bandwidth.
3. Provide a platform to develop and test the tracking algorithm code, and test with different processors.
4. Provide a platform to develop and test the trigger supervisor and monitor protocol. This includes initialization of the DSPs, fault detection, control and error messaging, and some hardware event histogramming.

The first two motherboards are assembled, have passed initial tests, and are providing the hardware platform for software development. The first DSP daughter card has been assembled, and initial tests have begun. This effort will continue through 2002.

7.7.4 Level 2

This section and the subsequent section present a progress report on the Level 2 and Level 3 triggers. Many of the details are ignored, since they have been presented in the BTeV proposal.

Since the BTeV proposal, the Level 2 and Level 3 code has been upgraded to handle the new pixel-detector geometry, which features z-staggered half-stations and a non-uniform checker-board of pixel chips. This was a relatively easy task because the Level 2 algorithm deals with three dimensional space points with coordinates that are expressed in a global coordinate system. Thus, only minor re-addressing schemes of the hit banks had to be done.

Unlike the Level 1 trigger R&D, we concentrate our effort for Level 2 on algorithm studies rather than CPU performance. Previous timing studies clearly indicate that our CPU and memory performance goals are achievable. Prior to describing the recent improvements, we now briefly describe the Level 2 trigger algorithm presented in our proposal. The input data

consists of (i) all Level 1 tracks and vertices (ii) the raw pixel hits (iii) optionally, hits from the first few forward tracking stations. Level 2 pixel tracks were seeded by Level 1 tracks. Missing hits were found, and these tracks were Kalman-fitted for the first time. In addition, new tracks of typically lower momentum were also reconstructed. Optionally, tracks were propagated forward and matching hits in the first few tracking stations were searched for, in order to improve the momentum resolution and provide “confirmation” for the tracks. Crude vertex fits were then performed. Detached tracks consistent with heavy quark decays surrounding the primary vertices were tallied. The final Level 2 trigger algorithm used detached vertices/tracks and the total detached p_t .

Based on further studies, we are now reaching the following conclusions:

- For forward tracking at Level 2, the advantages of using the first or second forward tracking stations is far from obvious. Although the first station does improve the momentum determination (from $\sigma_p/p \approx 5$ to 10% down $\approx 2\%$), such an improvement has only an indirect effect on the vertex. It simply allows us to measure the slope in the bend plane (y-z plane) a bit better. However, the y position resolution at the vertex does not improve significantly. Note that we currently have no direct momentum nor mass cuts in the vertex algorithm, to be able to trigger on various final states without bias. Although it has been proposed to confirm a pixel track with forward hits, the level of pattern recognition confusion in the first few tracking stations is much higher than in the pixel detector simply because the average detector occupancy jumps from $\approx 10^{-4}$ in the pixel to about 0.1 (first straw station).

Thus, we focus on improving the pattern recognition in the pixel detector rather than attempting to sort out partially reconstructed forward tracks. This could also simplify the data acquisition system and thereby improve the throughput, as we would not have to wait for the silicon strip or straw data. However, the data acquisition system is still designed to be able to deliver this data prior to a Level 2 accept, in case specific future Level 2 triggers are based on momentum or mass cuts rather than pure vertexing.

- Secondary vertex reconstruction at Level 2 improves substantially the signal to noise discrimination at Level 2. That is, after performing genuine vertex fits, we now require the presence of at least one fitted secondary vertex prior to computing the total detached p_t , which has to be greater than $2.5 \text{ GeV}/c$. In order to recover single charged track b vertices (for instance, $B^+ \rightarrow \pi^+\pi^0$), we alternatively require a single detached ($> 4\sigma$) track with $p_t > 2.0 \text{ GeV}/c$.

With these cuts, the Level 1 times Level 2 rejection rate on background crossings is about 1 in 1,000 and the Level 2 efficiencies for $B \rightarrow J/\psi K_s^0$ and $B^+ \rightarrow \pi^+\pi^0$ are about 90% and 94%, respectively. Histograms of the the total detached p_t are shown in Figure 7.40. The improvement relative to the results presented in the BTeV proposal comes mostly from the improved vertex fitting, while the pixel tracking performance remained virtually unchanged, despite significant changes in the design of the mechanical support and the geometry of the detector. Note that the single track efficiency for the Kalman fits could be improved upon,

as the description of the scattering surfaces in the fits could be more complete and precise. Also, a more sophisticated set of cuts will be possible in the future. Such cuts will be difficult to select without a more accurate description of the background, which is now dominated by beam crossings where multiple primary vertices occur within a few cm of each other.

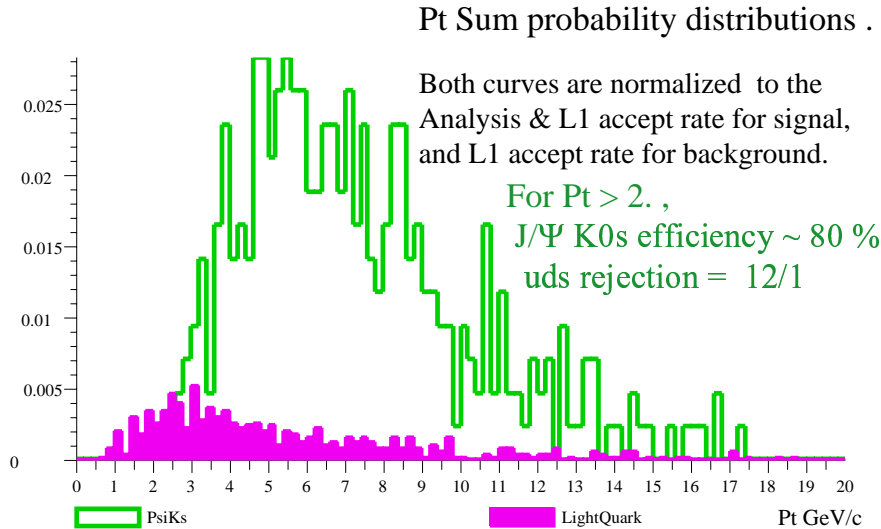


Figure 7.40: Distribution of the total detached transverse momentum for signal and background, for crossings with a secondary vertex. Note that an additional 10% efficiency comes from the second acceptance criteria (one single detached track with $p_t > 2.5$)

7.7.5 Level 3

Work on a prototype of the forward tracking reconstruction package has started. The goals are (i) to study and benchmark this critical component of all Level 3 triggers and (ii) optimize the design of these tracking elements. We are currently able to reconstruct tracks throughout the spectrometer, up to the EM calorimeter front face, without any prompting from the Monte-Carlo truth tables. Based on the same set of raw data structures and elementary algorithms, we are progressing on two distinct fronts: pixel-seeded tracks and K_s^0 's

7.7.5.1 Pixel to Forward Tracks

The reconstruction of the forward tracks coming from the interaction region is seeded by the pixel tracks (Level 2 tracks). Level 2 tracks are projected forward, to the z location of the straw and/or silicon strip stations. Note that we reconstruct tracks in both systems (straws

and silicon strip) in conjunction, as a typical track starts as silicon strip track but crosses into the straw system somewhere upstream of the RICH. For each track-station intercept, a track road (or foot-print) is computed based on the available fit. Hits are selected within these roads and a new Kalman fit is performed. For the straws, we fit a set of 2 or 3 hits within one stereo view to a given track rather than individual hits to avoid unnecessary combinatorics of many possible Kalman fits. Thus, left-right ambiguities are usually lifted prior to fitting. Arbitration among different possible Kalman fits for a unique Level 2 track seed is also performed at every tracking station. Such track-station Kalman-fitted intercepts are allowed to have one missing view. In addition, a single hopelessly confused tracking station can be skipped altogether.

The preliminary tracking efficiency versus momentum, for tracks reaching the RICH, is shown in figure 7.41. About half of the inefficiency is due to inaccuracies in the multiple-scattering accounting in the Kalman fits and the other half is due to pattern recognition confusion or double occupancy in the straws. The momentum resolution obtained via this full pattern recognition is in very good agreement with the fits performed in the context of BTeVGeant, where all hits are always assigned to their respective tracks. Once station 6 (located in front of the RICH) is reached, the probability of accepting “ghost” tracks is quite small, about 0.5%.

We anticipate improvement in the efficiency with a better version of the Kalman fitter. The pattern recognition in the first straw station is rather difficult, due to the high occupancy. We are considering increasing the size of the silicon strip station, and, possibly, entirely eliminating the straw planes at that location (95 cm downstream of the magnet center). Conversely, the silicon stations located just in front of the RICH or the EM calorimeter have negligible acceptance for pixel-seeded tracks.

7.7.5.2 Preliminary K_s^0 Tracking Studies

We now briefly describe the reconstruction of the K_s^0 for which we have no (or not enough) pixel hits. From a detailed study on the track topology of these $\pi^+ \pi^-$ pairs, we conclude that the largest reconstructable sample consists of tracks reaching straw station 6, for which we have 3 consecutive straw stations in the nearly field-free region beginning at $z \approx 2.75$ meters from the magnet center. Stations 4, 5 and 6 are located at $z \approx 2.9$, 3.3 and 3.8 meters from the magnet center, respectively. The following algorithm has been partly coded and is currently under study:

- Selection of “un-used” straw hits. We mark all the hits used in the above pixel-seeded Level3 tracks, as “used”, thereby getting rid of about 1/3 to 1/2 of the available hits.
- Reconstruction of straw hit triplets (or doublets) within a straw stack (or “view”). Despite the lack of good constraints from unknown track slopes, the multiplicity of such small, 2D tracks within a station and a stack is not overwhelmingly large.
- Reconstruction of 2D track between stations 4, 5 and 6, for each stereo view

Efficiency vs longitudinal momentum

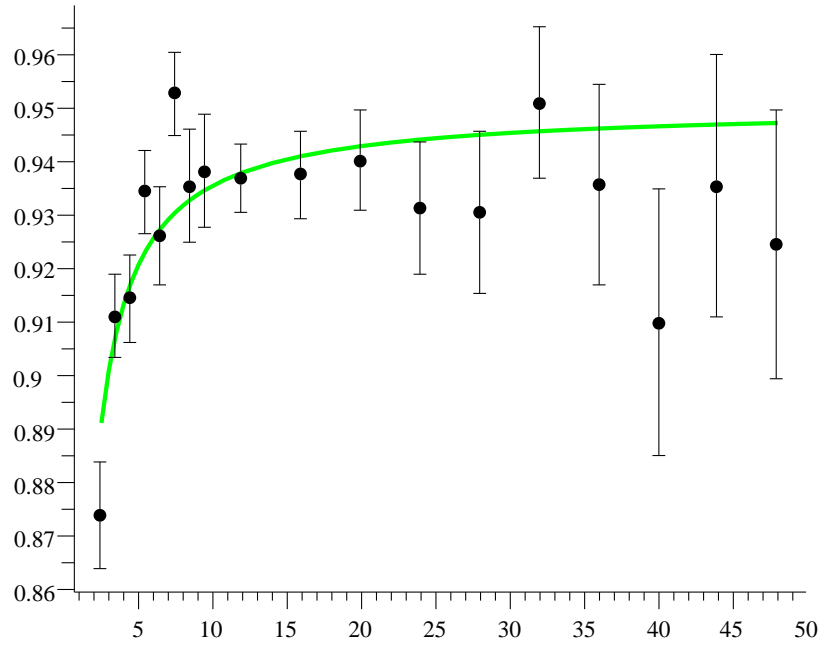


Figure 7.41: The preliminary efficiency versus the longitudinal momentum P_z (GeV/c) for the Level-2-seeded (pixel-seeded) Level 3 tracks. The fit is simply there to guide the eyes, the function is $E_\infty * (1 - k/P_z)$. The parameter k is obviously statistically significant, indicating problems at low momentum. Fortunately, this is where our geometrical acceptance drops sharply. Note that the loss of “efficiency” includes particles lost through interaction in the material of the detector and due to (recoverable) problems with the current detailed description of multiple scattering sources in the analysis program.

- 3D tracks between stations 4, 5, and 6.
- First reconstruction of a 3D K_s^0 vertex using the non-bend plane 2D vertex as a seed. The K_s^0 trajectory is constrained: it must come from the selected Level 2 primary vertex for which we have good detached p_t . This allows us to obtain a preliminary determination of the $\pi^+ \pi^-$ and K_s^0 momenta as well as the K_s^0 mass.
- Search for confirming hits in upstream stations, followed by track and vertex refits.

Preliminary studies indicate that we can reconstruct K_s^0 's that decay upstream of station 3 with about 60% efficiency. The loss of signal is mostly due to the high occupancy in the straws.

7.8 Data Acquisition R&D

7.8.1 Overview

The conceptual design of the BTeV readout and controls system has been outlined in the proposal. Since the proposal was submitted last year we have defined a baseline implementation, improved the overall throughput by 50% without increasing the costs and we began to evaluate the possible software architectures. This report begins with a brief system overview followed by more detailed discussion of the readout hardware and software as well as our plans for the detector control system.

By the time BTeV is operational the Tevatron will run with a bunch spacing of 132 ns corresponding to a crossing frequency of 7.6 MHz. For each crossing the BTeV detector will generate about 100 - 150 KBytes of data or roughly 1T Byte/second. The pixel detector dominates the data size with the second largest contribution coming from the electromagnetic calorimeter. Allowing for a 50% margin we designed the BTeV data acquisition system to handle a throughput of 1.5 T Bytes/s. Approximately 1% of the crossings will be accepted by the Level 1 trigger so that we have to route up to 15 GBytes/s to the Level 2/Level 3 trigger farm. Here the trigger rate will be reduced by a factor of 20 to approximately 4000 Hz. The expected reduction in data bandwidth is even larger as we plan to compress the data before it is sent to a mass storage system. In our baseline design we assume an output rate of 200 MBytes/s. A distributed software system will be developed to control and to configure the readout hardware.

Independent of the readout system, the BTeV detector control system will monitor the operation and performance of the many BTeV components. This system will be fully integrated with the Fermilab and Tevatron safety and security systems.

7.9 Data Acquisition Hardware

A schematic view of the readout system is shown in in the previous chapter. For each beam crossing the detector response is digitized and sent from the front-end electronics modules via serial links to so called Data Combiner Boards or DCBs. These custom designed modules combine data from several detector channels to an output stream which is sent over an optical fiber to the counting room. Here the data are buffered until a Level 1 trigger decision has been reached. Accepted events are passed on to the Level 2/Level 3 trigger farm and eventually transferred to a mass storage device. The design adopted for BTeV provides sufficient data throughput and minimizes the number of hardware modules that have to be developed by the collaboration. A single DCB design can easily be customized for different needs by different detector components. The data links and the Level 1 buffer modules are common to all of BTeV. Since the proposal was submitted we have refined the hardware architecture and decided to implement the system using multiple parallel data-paths called highways. Each highway (at least until a data fragment reaches the L2/L3 processing farm)

will process every n th crossing, where n is the total number of highways. The exact number of highways has yet to be determined, but is on the order of 8.

This approach greatly simplifies the data handling aspects of the system without compromising other requirements. For example, instead of a large number of small records a Level 1 buffer module now receives only $1/n$ th of the original number of messages. While the overall throughput remains the same - the average record is now larger - this is much better mapped to the performance characteristics of commercial network switches and hence easier to implement. A significant reduction in the number of control messages exchanged between the farm nodes, the Level 1 buffer system and the Global Level 1 processor(s) is another advantage of the highway architecture. During normal data taking operation the nodes of the Level 2/3 farm will be assigned to one of the highways. This will provide the highest throughput. All highways, however, will be cross connected via another network switch allowing each Level 2/3 computer access to the data in any of the Level 1 buffer modules. While the proposal kept open the option of custom designed switches we have now decided to use commercial Gigabit Ethernet switches.

In collaboration with the muon detector group we successfully tested a front-end board with ASDQ chips and an onboard serial data proving that with proper design digital noise does not affect the performance of nearby analog circuits.

Over the next year we intend to continue our tests and evaluations of components critical to the success of the BTeV data acquisition system. The activities planned include

- Development of a serial link board to test protocols, optical components and serializer/de-serializer chipsets.
- Evaluate FPGA chips. Implement a multi-channel TDC in an FPGA recently developed by Altera.
- Begin work on the Level 1 buffer design to provide a module for the trigger group allowing them to continue their development work.
- Development of the front-end part of the front-end - DCB link and integration on the next version of the muon readout board.
- Complete requirement documents and specifications.

7.9.1 Data Acquisition Software

A vast amount of software is needed to operate an experiment of BTeV's complexity. Broadly speaking this can be divided into the following categories

- Run Control including configuration of the readout hardware
- Detector and environmental control
- Error and alarm handling and recovery

- User interface and remote access including security
- Persistency and archive

Since the submission of the proposal we have continued to develop the software architecture of the BTeV data acquisition system and defined the requirements and functionality of the software components listed above. This process will continue with the expectation to have a detailed model of the readout and control software by the end of the year. In parallel to this effort we have chosen two critical components for a careful evaluation of commercial and free solutions. The first is the message passing system, which provides the communication path between all software components. The proper selection of this tool is of critical importance. The other component we will evaluate over the next year are commercial detector/process control systems. We will try to learn which system is best suited for the control and monitoring of the BTeV detector and its environment. We will also investigate a recent trend in the HEP computing community to use similar tools for detector and run control and to establish an overall experiment control system. As with the readout hardware we will base our system on commercial components and take advantage of successful software packages developed by the HEP community. We are establishing contacts with other laboratories and in particular with those groups involved in data acquisition software development for LHC experiments.

7.9.2 Summary and Future Plans

The conceptual design of the BTeV readout system is well advanced. For our baseline implementation we have relied on commercial products and industrial standards wherever possible. Interface modules and data links to the front-end modules as well as the intermediate event buffers have to be custom designed but no new technologies have to be developed for these purposes. Over the next year we will begin to evaluate commercial network components and the design of critical components for the data combiner and the Level 1 buffer modules will start as well. On the software side we plan to evaluate commercial and freeware message passing systems that will form the core layer of our distributed run control system. On top of that we will complete the design of the system architecture and define the interfaces between different components so that development work can proceed at several locations. As part of this we will further develop the interface to the RTES effort. We intend to base the detector control system on commercial software packages. Initial contacts with vendors have been established. Within the next year the evaluation phase should be completed.

7.10 BTeV Real-Time Embedded Systems R&D

7.10.1 Introduction

A major challenge facing BTeV is the software infrastructure required to keep the trigger system operating, to assure that it is working correctly, and to detect and adapt to fault conditions both within the trigger itself and within the experiment and accelerator environment. BTeV physicists have formed an alliance with computer scientists and engineers at four universities (Illinois, Pittsburgh, Syracuse, and Vanderbilt) in part to address these issues. This collaboration is researching the design and implementation of high-performance, heterogenous, fault-tolerant and fault-adaptive real-time embedded systems, of which the BTeV trigger is an excellent example. This research is being funded by a \$5 million, 5-year award from the National Science Foundation's Information Technology Research program. While this research will be applicable to a wide variety of problems in science, medicine, and industry, one important deliverable of the research is an operating system for the BTeV trigger.

7.10.2 Project Overview

The BTeV trigger system is an example of a very large-scale real-time embedded computer system that:

- achieves ultra high computational performance through use of parallel hardware architectures;
- achieves and maintains functional integrity via distributed, hierarchical monitoring and control;
- is required to be highly available; and
- is dynamically reconfigurable, maintainable, and evolvable.

BTeV will produce very large streams of data which must be processed in real-time using data dependent computation strategies. Such systems are inextricably tied to the environment in which they must operate, and must perform complex computations within the timing constraints mandated by their environments. These systems require ultra high performance (on the order of 10^{12} operations per second). This level of performance requires parallel hardware architectures, which in the case of BTeV is composed of a mix of thousands of commodity processors, special purpose processors such as Digital Signal Processors (DSPs), and specialized hardware such as Field Programmable Gate Arrays (FPGAs), all connected by very high-speed networks. The systems must be dynamically reconfigurable, to allow a maximum amount of performance to be delivered from the available and potentially changing resources. The systems must be highly available, since the environments produce the data

streams continuously over a long period of time, and interesting phenomena important to the analysis being done are rare and could occur in the data at any time. To achieve the high availability, the systems must be fault tolerant, self-aware, and fault adaptive, since any malfunction of processing elements, the interconnection switches, or the front-end sensors (which provide the input stream) can result in unrecoverable loss of data. Faults must be corrected in the shortest possible time, and corrected semi-autonomously (i.e. with as little human intervention as possible). Hence distributed and hierarchical monitoring and control are vital.

The design and implementation of such systems cannot be achieved by the ad hoc approach of developing simple small-scale components and scaling them up into large-scale systems. Creating usable software for this type of real-time embedded system will require research into solutions of general problems in the fields of computer science and engineering. We plan to approach these problems in a way that is general, and to produce methodologies and tools that can be applied to many scientific and commercial problems. Issues such as fault tolerance and performance must be explicitly addressed at multiple levels in the system design. We propose advances in system design methodology, tools and runtime infrastructure to facilitate these and more issues involved in developing such systems. We further propose to develop the software to accomplish the design and implementation of the system and to study its performance, utility, and scalability on the actual BTeV hardware as it grows over the construction phase of the experiment. The result of this research will be software, design methodologies, and the documented experience of the project.

Several capabilities are required:

1. **System Modeling and Analysis** - Full-system performance estimations are needed during development, given the coupling that exists between different aspects of a system design (e.g. low-level architectural decisions can have a large impact on system-level performance and fault behavior). Designers need mechanisms for representing and evaluating the impact of these decisions. The design tool will serve as a framework for modeling and analyzing system designs via behavioral simulation, performance simulation, design verification, etc. The design tool will continue to be useful during operations to understand how to handle unanticipated situations, which often arise in HEP research;
2. **System Configuration Management** - Configuration of a large-scale networked processing system is a complex problem, more so when the system is susceptible to faults. A robust configuration management infrastructure is required, with the ability to specify reconfiguration strategies at different levels. The fault mitigation infrastructure is intricately coupled to the configuration management infrastructure - means to capture the specifics of the coupling are necessary;
3. **Runtime Environment and Hierarchical Fault Detection/Management** - The deployment, execution and reconfiguration of the components must be carefully managed, especially when the cost of downtime is high, as is the case for BTeV. Runtime

environment control is essential. A system-wide infrastructure is required for rapidly detecting, isolating, filtering, and reporting faults. In very large-scale heterogeneous systems a single centralized fault management solution is clearly not feasible. Hierarchical distributed fault mitigation is necessary, with the ability to specify fault mitigation policies at different levels of abstraction (system, network, node, etc.).

The researchers on this project have extensive experience in developing all of the above capabilities. For example, a high-level design tool is required to support the overall design, deployment, and evolution of BTeV-type systems. The Model Integrated Computing (MIC) approach, developed at the Institute of Software Integrated Systems (ISIS), Vanderbilt University, assists the creation of domain-specific modeling, analysis, and program synthesis environments for building complex, large-scale computer-based systems. Integrated models created in this environment represent all relevant factors of a physical system. Models can be subjected to many types of rigorous analysis, for verifying the behavior and performance of the system prior to implementation. Central to this approach is the concept of a “configuration,” which is a particular organization of computing resources, such as processors, network components, memory buffers, and storage elements, and a particular allocation of software components and datasets on them, including task schedules and message routes. Systems can be synthesized (or generated) from the models when the designer is satisfied with the analysis results. The Design and Analysis Environment will consist of: a) graphical modeling language/environment for system specification; b) synthesis tools for interfacing the models with commercially available and custom analysis tools; and c) synthesis tools for generating configurations from specifications, for configuring fault managers, and for configuring operation managers. The synthesized configurations are deployed and executed in the Runtime Environment. The primary interaction between the Design and Runtime environments is through the synthesis process. However, feedback from the Runtime to the Design environment is possible in advanced fault scenarios that require re-synthesis and re-deployment. While MIC provides the basic infrastructure, research is required to define: (1) modeling language and composition methodologies suitable to BTeV’s application; (2) mapping techniques for models to/from analysis tools; (3) large-scale synthesis techniques.

The Illinois, Pittsburgh, and Syracuse groups have extensive experience in fault detection and mitigation, as well as real-time operating systems. Very Light Agents (VLAs), developed by the groups at Pittsburgh and Syracuse, will be applied at the lowest level of the runtime hierarchy. Applied to DSPs, these are simple software entities, which can be implemented in a few dozen lines of assembly language, that take advantage of the exception-signaling and interrupt-handling mechanisms present in most DSP kernels to expose errors in the kernel behavior. When the VLA detects (e.g., by monitoring DSP exception signals) an error condition, it will report to an ARMOR (described next), which will take appropriate action such as disabling the execution thread or discarding the current data item. A similar mechanism will be explored for the monitoring and reporting of deadlines, traffic, processor loads, etc. Moreover, the interrupt mechanism will also be used to trigger reconfiguration of the software or hardware at this lowest level of the hierarchy. Note that since the software VLAs are small and interrupt-driven, the latency introduced by VLAs will be negligible.

Hardware VLAs can also be developed for FPGAs, consuming only small number of gates, and taking advantage of otherwise present communication resources. VLAs (software and hardware) in this context is new research area.

The fault tolerance and performance-oriented services offered to the system will be encapsulated in intelligent active entities (agents) developed at the University of Illinois called ARMORs (Adaptive, Reconfigurable, and Mobile Objects for Reliability). ARMORs are, by design, highly flexible processes, which can be customized to meet the runtime needs of the system. Variants of ARMORs will run on DSPs, L2/L3 processors, and other supporting processors throughout the system. ARMORs communicate through message passing and all functions of an ARMOR process and its runtime behavior are encapsulated in “elements.” Elements constitute basic building blocks, which usually encapsulate elementary detection and recovery services available to the application. New functionality can be introduced into the system without disturbing existing functionality, as long as, from a resource or timing perspective, it does not affect the current system. In other words, the resource manager must implement some type of resource protection. Services provided by the elements are invoked by the ARMOR interface, which serves as a communication gateway with the outside world. The ARMOR interface has two primary responsibilities: (1) controlling the addition, removal, and replacement of constituent elements within the ARMOR, and (2) providing communication among ARMORs. An application can take advantage of ARMOR provided services (such as error detection and recovery) through the concept of an embedded ARMOR in which the core element structure of the ARMOR is linked to the user application process. The application code is lightly instrumented with the embedded ARMOR API to invoke the services provided by the underlying elements. In this configuration, the embedded ARMOR process appears as a full-fledged ARMOR to other ARMORs in the system and as a native application process to non-ARMOR processes. This permits BTeV physics applications to use the same apparatus for error handling as will be used to handle errors within the computing platform itself.

7.10.3 Strong Connection to BTeV

We believe that there are very significant advantages to connecting this research to the BTeV experiment. Not only will the software and methods produced by this research have significant impact on one of the most important areas of investigation in HEP, but the generalizable computer engineering research will also be directly applicable to a large class of similar real-time embedded computer systems. The BTeV trigger system hardware, which will be provided by Fermilab as part of the experiment, will supply an extremely important ingredient in this project: a large test-bed that represents millions of dollars of equipment and comes with a highly motivated set of users who will test the methodologies and tools developed in an extremely harsh environment over an extended period of time. The test-bed will be built gradually as the proposed research progresses, from a 5% system in 2002 to a full system in 2006-2007. It will therefore be possible for the software developers, aided and supported by the experimenters, to test and refine the software and strategies continuously

and incrementally throughout the lifetime of this project. The close interdisciplinary contact between the experimenters and computer scientists will also help introduce important computer science research into the HEP community, which has not always been aware of work that has been done in this area and has not taken full advantage of it.

7.10.4 Group Members

The team that has been assembled to carry out this research consists of the leaders of the BTeV trigger and data acquisition system development efforts and Computer Scientists/Engineers who are experts in the field of embedded systems, real-time systems, and fault tolerant computing. The Computer Scientists/Engineers come from the University of Illinois, the University of Pittsburgh, Syracuse University, Vanderbilt University, and Fermilab. The team is committed to carrying out the proposed R&D and implementing a series of systems of increasing size and complexity, using the experience gained at each stage to refine and improve the system until it is demonstrated to scale to the full BTeV system. More information on the research groups involved in this project is available at the project's website (http://www.hep.vanderbilt.edu/btev_rtes/).

7.10.5 Current and Proposed Activities

The group has been meeting regularly for 6 months, although the funding did not become available until the end of 2001. The computer science/engineering groups are just beginning to hire students and post-docs and are ramping up their operations. The group has written a work plan for their first year activities, and produced milestones for all five years of the project. These milestones are listed in Table 1.

Table 7.4: Project Milestones

Funding Year	Design Environment Milestone	Runtime System Milestone
FY1: Q1–2	Modeling language and environment (preliminary). Specify Interface to runtime environment.	Design of overall runtime system hierarchy (ARMOR + VLAs).
FY1: Q3–4	Synthesis of operations and Fault managers. DSP and LINUX Synthesis.	Design and implementation of VLA & ARMOR prototypes
FY2: Q1–2	Modeling language and environment. Design space (preliminary).	Communication structure between VLAs and the levels above.
FY2: Q3–4	Synthesis of performance simulator. Synthesis of all operations managers (final). Hardware synthesis.	Detection and recovery in Layer 1 of ARMOR. Study Dynamic load-balancing (DL).
FY3: Q1–2	Modeling language and environment (final). Design space.	Detection and recovery in Layer 2 of ARMOR; Study DL.
FY3: Q3–4	Synthesis to Diagnosability tool. Synthesis to performance simulator (final).	Detection and recovery in Layer 3 of ARMOR; Study DL.
FY4: Q1–2	Design space (final).	Full scale Runtime Environment test.
FY4: Q3–4	Synthesis to Reliability tool. Synthesis to Diagnosability tool (final).	Large scale evaluation on BTeV hardware and revision.
FY5: Q1–2	Synthesis to Reliability tool (final).	Final evaluation on BTeV hardware.